EXHIBIT D

IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF MASSACHUSETTS

SKYLINE SOFTWARE SYSTEMS, INC., Plaintiff, v. KEYHOLE, INC., and GOOGLE INC. Defendants.	CIVIL ACTION NO. 10980-DPW

EXPERT REPORT OF PROFESSOR STEVEN K. FEINER, Ph.D.

December 8, 2006

Nelson G. Apjohn (BBO No. 020373)
NUTTER McCLENNEN & FISH LLP
World Trade Center West
155 Seaport Boulevard
Boston, MA 02210
(617) 439-2000
Fax: (617) 310-9000

Attorneys for Defendants and
Counterclaimants
KEYHOLE, INC. and GOOGLE INC.

<u>Of Counsel</u>
Darryl M. Woo, admitted *pro hac vice*
Michael J. Sacksteder, admitted *pro hac vice*
Heather N. Mewes, admitted *pro hac vice*
Carolyn Chang, admitted *pro hac vice*
Saundra L. M. Riley, admitted *pro hac vice*
FENWICK & WEST LLP
Embarcadero Center West
275 Battery Street
San Francisco, CA 94111
Tel. (415) 875-2300
Fax (415) 281-1350

I. INTRODUCTION

1. This report addresses the issue of whether U.S. Patent No. 6,496,189 ("the '189 patent"), directed to remote landscape display and pilot training, is or is not valid. I have been asked by defendants Keyhole, Inc. and Google, Inc. (collectively, "Google") to provide at this stage of the case my expert opinion regarding whether or not claims 2-3, 5, 7-9, 11, 13-14, 16, 18-19, and 21-24 of the '189 patent are valid in light of certain prior art references and to supplement my reports of July 10, 2006, and July 29, 2006, regarding claims 1 and 12 in order to address the Court's new claim construction order as well as subsequent discovery. I reserve the right to supplement this report to address the validity or invalidity of any other claims of the '189 patent and to address any additional claim construction or discovery relating to validity.

2. As discussed in detail below, in my opinion, claims 1-3, 5, 7-9, 11-14, 16, 18-19, and 21-24 of the '189 patent are invalid. Each of these claims is anticipated and/or rendered obvious by the following prior art references: (1) the TerraVision system; (2) the T_Vision system; and (3) Migdal and Cosman.

II. QUALIFICATIONS AND PUBLICATIONS

3. My qualifications are stated more fully in my *curriculum vitae*, a true and correct copy of which is attached as Exhibit A. However, I will provide a brief summary of my relevant qualifications here.

4. I received a Ph.D. in Computer Science from Brown University in 1987. I received an A.B. degree in music from Brown University in 1973.

5. I am presently a Professor of Computer Science at Columbia University, a position that I have held for over twenty years. I have been a Full Professor since January of 2000. Prior to that, I was an Associate Professor of Computer Science at Columbia University from January of 1991 until December of 1999, and an Assistant Professor from September of 1985 to December of 1990. Prior to joining the faculty of Columbia University in September of 1985, I was a Research and Teaching Assistant in the Department of Computer Science at Brown University from September of 1977 until August of 1985.

6. At Columbia University, I direct the Columbia University Computer Graphics and User Interfaces Laboratory, and teach both graduate and undergraduate students in computer graphics and user interfaces courses. I advise Computer Science doctoral candidates, primarily in the field of computer graphics and user interfaces. I am an active academic researcher, whose areas of research include knowledge-based design of graphics and multimedia, user interfaces, virtual reality and augmented reality, wearable computing, animation, hypermedia, and visualization.

7. I am a coauthor of *Computer Graphics: Principles and Practice*, *Second Edition*, Addison-Wesley, 1990 ("*Computer Graphics*"), an authoritative and frequently cited academic computer graphics text. I am also a coauthor of *Introduction to Computer Graphics*, Addison-Wesley, 1993, and *Computer Graphics: Principles and Practice, Second Edition in C*, Addison-Wesley, 1996. As indicated on my curriculum vitae, I am the author or a coauthor of over thirty journal papers, over seventy conference papers, and numerous other workshop papers, books and

book chapters, editorials and other publications on computer graphics and user interfaces. I have been an Associate Editor of *ACM Transactions on Graphics* and *ACM Transactions on Information Systems*, and have been on the editorial boards of *IEEE Transactions on Visualization and Computer Graphics*, and *Virtual Reality*. I am a frequent invited speaker on computer graphics and user interfaces at institutions such as Princeton University, the Massachusetts Institute of Technology, and Carnegie Mellon University. In addition, I have given invited talks at numerous conferences and workshops, including ones related to Geographic Information Systems ("GIS"), such as *GIScience 2002*, the Advanced Research and Development Activity *Geospatial Intelligence Information Visualization Researchers Meeting 2003* and *GIS Planet 2005*. In 1991, I received an Office of Naval Research Young Investigator Award.

8. I am a named inventor on an issued United States patent relating to computer graphics entitled "Worlds-within-worlds nested display and interaction system and method" (U.S. Pat. No. 5,524,187).

III. DATA AND INFORMATION CONSIDERED

9. A complete list of the data and other information I considered in forming the opinions stated in this Expert Report is attached as Exhibit B.

10. I have reviewed the '189 patent and its relevant prosecution history and am familiar with this patent, its claims, and the background technology. A copy of the '189 patent is attached hereto as Exhibit C. I am also familiar with the Court's claim construction orders dated March 24, 2006 and November 16, 2006, attached hereto as Exhibits D-1 and D-2.

11. I have also reviewed certain prior art references relating to the TerraVision system, the T_Vision system, and Migdal and Cosman. These include, but are not limited to, the following:

12. For the TerraVision system:

a. Y.G. Leclerc, "MAGIC Final Report," SRI International, Menlo Park, CA (May 1996), available at http://www.ai.sri.com/~magic/magic-final-report.html (GOOG 000358-70 & Lau Depo. Ex. 86), attached hereto as Exhibit E ("MAGIC Final Report").

b. Y.G. Leclerc & S.Q. Lau, Jr., "TerraVision: A Terrain Visualization System," Technical Note 540, SRI International, Menlo Park, CA (April 22, 1994), available at http://www.ai.sri.com/pubs/files/778.pdf (GOOG 000371-390 & Lau Depo. Ex. 87), attached hereto as Exhibit F ("TerraVision Technical Note").

c. Barbara Fuller & Ira Richer, "The MAGIC Project: From Vision to Reality," *IEEE Network*, Vol. 10, No. 3, pp. 15-25 (May/June 1996) (GOOG 000346-25 & Lau Depo. Ex. 88), attached hereto as Exhibit G ("MAGIC IEEE Article").

d. TerraVision: A High Speed Terrain Visualization System (1994) and Architecture and Initial Performance of TerraVision (1994) (G-T_0018 & Lau Depo. Ex. 202) ("TerraVision Videos"). A transcript of the TerraVision Videos was made during the Deposition of Stephen Lau (Lau Depo. at 164:19-167:20, 171:6-174:24), and is attached hereto as Exhibit H.

e. TerraVision Source Code (April 1996) (G-T_0020 and Lau Depo. Exs. 206-220).

13. For the T_Vision system:

a. U.S. Pat. No. 6,100,897, filed Dec. 17, 1996 (GOOG 009064-81), attached hereto as Exhibit I ("Mayer patent").

b. Counterpart German application DE 195 49 306 A1, filed Dec. 22, 1995 (GOOG 0021372-88 (translation)), attached hereto as Exhibit J ("Mayer German patent application").

c. *SIGGRAPH '95 Visual Proceedings* (Lau Depo. Ex. 205 at GOOG 0017566-69, 0017699), attached hereto as Exhibit K.

d. SIGGRAPH '95 Multimedia CD-ROMs, \COMUNITY\TVISION (G-T 0021) ("T_Vision Project"). Printouts of relevant pages from these CDs are attached hereto as Exhibit L.

e. Terra1995 Video (G-T_0013) ("T_Vision Video"). A transcript of the T_Vision Video (GOOG 27815-16) is attached hereto as Exhibit M.

14. For Migdal and Cosman:

a. U.S. Pat. No. 5,760,783, filed Nov. 6, 1995 (GOOG 000290-315), attached hereto as Exhibit N ("Migdal patent").

b. Counterpart International Application PCT/US96/17673, filed Nov. 6, 1995 (GOOG 0021389-448), attached hereto as Exhibit O ("Migdal International Application").

c. Michael Cosman, "Global Terrain Texture: Lowering the Cost," *Proceedings of the 1994 Image VII Conference*, Tempe, Arizona, The IMAGE Society, pp. 53-64 (GOOG 000334-45), attached hereto as Exhibit P ("Cosman").

15. In addition, the following prior art references were relevant to my invalidity analysis:

a. Boris Rabinovich and Craig Gotsman, "Visualization of Large Terrains in Resource-Limited Computing Environments," *Proceedings of*

Visualization '97, Phoenix, Arizona, pp. 95-102 (GOOG 9826-35), attached hereto as Exhibit Q ("Rabinovich").

b. Michael Potmesil, "Maps alive: viewing geospatial information on the WWW," *Computer Networks and ISDN Systems*, Vol. 29, Nos. 8-13, pp. 1327-42 (Sept. 1997) (GOOG 009760-77), attached hereto as Exhibit R ("Potmesil").

c. United States Patent No. 5,828,844, filed October 8, 1996 (GOOG 030235-44), attached hereto as Exhibit S ("Civanlar").

d. Internet Society Network Working Group, Request for Comments: 1577 (January 1994) (GOOG 030196-211), attached hereto as Exhibit T ("RFC 1577").

e. Internet Society Network Working Group, Request for Comments: 1821 (August 1995) (GOOG 030212-34), attached hereto as Exhibit U ("RFC 1821").

16. I have also reviewed the depositions of Stephen Lau, Ronnie Yaron, Ofer Shor, and Michael Jones.

IV. OPINIONS AND BASES/REASONS FOR OPINIONS

17. I have been asked to provide my opinion regarding whether or not certain prior art references invalidate the asserted claims of the '189 patent, namely claims 1-3, 5, 7-9, 11-14, 16, 18-19, and 21-24.

18. It is my understanding that a claim is invalid as anticipated if every limitation of that claim is disclosed in a single prior art reference either expressly or inherently. It is my further understanding that a claim is invalid as obvious if differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art. I also understand that to prevent hindsight invalidation of patent claims, the law requires some teaching or motivation to combine the prior art references.

A. Tutorial Regarding Relevant Technology

19. If requested, I may provide a tutorial explaining the state of the computer graphics art prior to the filing of the '189 patent, including concepts that were well-known and understood in the field at that time. I may also reference and explain the prior art described and cited in the '189 patent, including the "references cited" on the cover of the patent, the prior art systems and methods discussed in the section labeled "background of the invention," and the "references cited" on the cover of U.S. Patent No. 6,111,583, which is incorporated by reference in the '189 patent. In addition, I may explain the disclosures in the specification of the '189 patent and describe the claims of that patent as construed by the Court.

B. <u>The Level of Ordinary Skill in the Art</u>

20. I understand that patent claims should be construed from the perspective of a person of ordinary skill in the relevant art at the time the patent application was filed. In my opinion, the relevant art for the '189 patent is computer graphics. The patent-in-suit uses concepts, nomenclature, designs, and systems from the computer graphics art that should be understood in this context. In my opinion, one of ordinary skill in the art relevant to the subject matter of the '189 patent at the time the application for the patent was filed would be a person with a bachelor's degree in Computer Science including at least one course in computer graphics, or with academic or work experience equivalent to that level of education.

C. <u>The TerraVision System</u>

21. In my opinion, the TerraVision system itself, as well as certain publications describing this system, anticipate and/or render obvious claims 1-3, 5, 7-9, 11-14, 16, 18-19, and 21-24 of the '189 patent. These publications include the MAGIC Final Report, the TerraVision Technical Note and the MAGIC IEEE Article.

1. *Overview of the TerraVision System*

22. TerraVision was a system for visualization of terrain. It allowed a user to view, in real time, a synthetic recreation of a real landscape created from elevation data and a large number of aerial or satellite images of that landscape. The goal of the system was to allow a user to roam about the terrain at will.

23. Like the '189 patent, the TerraVision system included the following capabilities:

- "to display on a remote computer earth image data that was stored on a server somewhere else … from a viewpoint chosen interactively by the user"

- "to train a pilot to fly a preplanned flight course while allowing the pilot to see the view, see it at any point along the flight course at substantially any desired angle"

- "to display on a client computer images of ground terrain stored in a remote server which are conveyed to the client via a network [or] via the Internet"

- "to stream data required for rendering three-dimensional terrain images on a remote computer"

[Lau Depo. at 51:8-53:18.]

24. It is my understanding, based on the documents and testimony I have reviewed regarding TerraVision, that SRI International, and specifically Yvan Leclerc and Stephen Lau, developed the TerraVision system in the early to mid 1990s. SRI is an independent non-profit research institute. It developed TerraVision as part of a government-funded research project called MAGIC ("Multidimensional Applications and Gigabit Internet Consortium"). The MAGIC project had three components: a high-speed network, a network-based storage system, and TerraVision. SRI was responsible for developing TerraVision.

25. The TerraVision system was conceived and reduced to practice in the United States earlier than May 1996. Stephen Lau testified that the system was conceived in 1992 and reduced to practice by late 1993 or early 1994:

Q. Could you describe, please, what TerraVision was?

A. TerraVision was an application that allowed the user to manipulate or fly through, do aerial fly throughs of terrain in real time over a network or the Internet, and what you would see of what were all the imagery was potentially located on a separate system or separate server. As depended upon the viewpoint of the person, TerraVision would request from servers spread out on a network in arbitrary locations imagery of different resolutions based upon the location of the viewer where higher-resolution imagery, you would want higher-resolution imagery for images close to the user and lower resolution far off in the distance, in the horizon. It had a 2D component where you could pan and zoom and also the three-dimensional component where you could actually do a real-time fly through. Later on in the project it was expanded. The application was expanded such that you could do it in stereoscopic views -- stereographic views.

…

Q. At about what time were all these features included in the TerraVision software?

A. The features were in place initially during the design phase which happened around ’92 or so. By 1994, early 1994, all the components were in there. Actually, it’s probably a little earlier than that. Probably late ’93.

Q. So by all components, you’re referring to the things you’ve mentioned in your answer a few minutes ago, the ability to fly through terrain in real time over the network, seeing the images from a separate server and so forth?

A. Yes.

[Lau Depo. at 32:25-33:21, 34:6-21 (objection omitted).] This is corroborated by overwhelming evidence. The TerraVision system is described in detail in an SRI Technical Note, approved for public release by SRI no later than January 26, 1995. [*See* Ex. F; *see also* Lau Depo. at 95:10-103:16.] It is also described in the MAGIC Final Report submitted to the Defense Advanced Research Projects Agency (“DARPA”) in May 1996. [*See* Ex. E; *see also* Lau Depo. at 57:12-58:6, 59:20-60:14, 62:2-82:9.] Both of these publications were made available on SRI’s website for the MAGIC project. [*See, e.g.,* GOOG 24834-40 (pages from SRI website dated October 1996 showing links to both the MAGIC Final Report and the TerraVision Technical Note).] I have confirmed this by retrieving and viewing archived versions of the SRI website dating from as early as May 1997 and October 1997 using the “Wayback Machine” from the Internet Archive (http://www.archive.org). Furthermore, numerous articles and other contemporaneous

publications refer to the TerraVision system. [*See, e.g.*, Agranov, G. and Gotsman, C., "Algorithms for rendering realistic terrain image sequences and their parallel implementation," *The Visual Computer*, vol. 11, pp. 455-464, 1995, obtained from http://www.cs.technion.ac.il/~gotsman/AmendedPubl/AlgorithmsForRender/AlgorithmsForRender.pdf; Tierney, B., Johnston, W., Herzog, H., Hoo, G., Jin, G., Lee, J., Chen, L., and Rotem, D., "Distributed Parallel Data Storage Systems: A Scalable Approach to High Speed Image Servers," *Proceedings of the Second ACM International Conference on Multimedia*, San Francisco, CA, 1994, pp. 399–405, obtained from http://doi.acm.org/10.1145/192593.192709.] An article published in *IEEE Network* in May/June 1996 provides yet another detailed description of the system. [Ex. G; *see also* Lau Depo. at 103:17-108:10.] I have also reviewed quarterly reports on the MAGIC project submitted by SRI to DARPA for the period January 1993 to July 1995, and another SRI proposal to DARPA from June 1996 describing the TerraVision system. [*See* Lau Depo. Exs. 90-99, 200 (SRI Quarterly Reports covering the periods Q4 1992 to Q2 1995); Yvan G. Leclerc et al., "BADD Data Aggregation and Visualization," Technical Proposal, SRI International, Menlo Park, CA (June 17, 1996); *see also* Lau Depo. at 110:11-147:14.] These reports detail SRI's work on TerraVision. In addition, Mr. Lau has provided his engineering notebooks, reflecting his work on TerraVision. [Lau Depo. Exs. 223, 225; *see also* Lau Depo. at 229:3-230:19, 232:25-234:4.] The TerraVision system was not abandoned, suppressed or concealed; indeed, the evidence is overwhelmingly to the contrary. The TerraVision system was publicly demonstrated, and information about this system was widely disseminated, as detailed herein.

26. Stephen Lau has also provided source code for the TerraVision system. [*See* Lau Depo. at 216:6-224:25 & Lau Depo. Exs. 206-20.] The TerraVision source code describes the operation of the TerraVision system. I have reviewed the revision history for the source code, as reflected in the code itself, and this version of the source code dates from no later than April 27, 1996. In fact, a number of the files were last revised even earlier. I have also confirmed with Stephen Lau that he obtained this version of the source code prior to leaving SRI International in approximately May 1996. He maintained the source code on a server at Lawrence Berkeley National Labs ("LBNL") until approximately March 2006, when he went to work for the University of California. This is corroborated by the "date modified" for many of the source code files, which is listed as February 21, 2006. This would be the date the files were copied from the LBNL server to a CD. There is no indication that any of the source code itself was modified or otherwise altered after April 27, 1996.

27. The TerraVision system was also in public use and publicly known and used by SRI and others in the United States earlier than May 1996. This is confirmed by the evidence cited above. Furthermore, it is corroborated by evidence that TerraVision was publicly demonstrated at a number of public conferences and symposia prior to May 1996, including SIGGRAPH '95 (*see* Lau Depo. at 82:10-84:4, 193:16-194:18, 201:5-202:11, 206:2-214:24 & Lau Depo. Ex. 205; GOOG 24714-17 (pages from the SIGGRAPH '95 website available at http://www.siggraph.org/s95/)), Supercomputing '94 and Supercomputing '95 (*see* Lau Depo. at 84:5-86:14), the 1994 MAGIC Symposium (*see* Lau Depo. at 159:11- 162:23, 169:1-12, 175:6-8 & Lau Depo. Ex. 201), and the 1995 Magic Symposium (*see* Lau Depo. at 88:15-90:6, 93:20-95:3, 195:2-198:2 & Lau Depo. Ex. 204). [*See also* Ex. E at GOOG 366-70; Ex. G at GOOG 355; Lau Depo. at 198:3-200:6 (describing other public demonstrations).] The system demonstrated at SIGGRAPH '95 had (1) a local memory which stored data blocks corresponding

to coordinates proximal to a current viewpoint of the renderer; (2) a communication link through which the memory received the data blocks from a remote server; and (3) a processor which received one or more specified coordinates along with indication of a respective resolution level from a renderer, provided the renderer with a first data block which includes data corresponding to the one or more specified coordinates from a local memory, and downloaded over the communication link one or more data blocks of a resolution level higher than the resolution level of the first block which include data corresponding to the one or more coordinates if the first block is not from the indicated level. [Lau Depo. at 206:8-207:2 (claim 12); *see also* 193:16-21 (claim 1).] I also attended SIGGRAPH '95 in Los Angeles, California, and recall seeing TerraVision being demonstrated in the Interactive Communities venue. That TerraVision was in public use earlier than May 1996 is also corroborated by evidence of two demonstration TerraVision Videos shown at the 1994 MAGIC Technical Symposium, and made available to the public by SRI. [Ex. H; *see also* Lau Depo. at 164:1-178:11, Lau Depo. Ex. 97 at GOOG 26595 (Q3 1994 Quarterly Report) & Lau Depo. Ex. 201 at GOOG 26979-84 (script for videos from 1994 MAGIC Technical Symposium proceedings).]

28. The MAGIC Final Report (Ex. E), the TerraVision Technical Note (Ex. F), and the MAGIC IEEE Article (Ex. G) are themselves printed publications, published more than one year before the filing date of the '189 patent.

29. The application for the '189 patent was filed on February 26, 1999. [*See* Ex. C.] An earlier related application was filed on September 29, 1997, which issued as US Patent No. 6,111,583 ("the '583 patent"). Thus, it is my understanding that the earliest priority date that Skyline could possibly claim for the '189 patent is September 29, 1997. The TerraVision system was in public use more than one year prior to this date, and the MAGIC Final Report (Ex. E), the TerraVision Technical Note (Ex. F), and the MAGIC IEEE Article (Ex. G) were all published more than one year prior to this date.

30. Furthermore, it is my opinion that the '189 patent is not entitled to the September 29, 1997 priority date of the '583 patent, because not a single claim of the '189 patent is supported by the specification of the '583 patent. Every claim of the '189 patent involves "downloading from a remote server" or a processor that "downloads over the communication link" from "a remote server". However, the specification of the '583 patent does not disclose the existence of a remote server, let alone downloading from one. In addition, I understand that Skyline has admitted that it is not asserting the September 29, 1997 priority date. [*See* Skyline Response to RFA No. 1.]

31. I am also familiar with the deposition testimony of Ronnie Yaron and Ofer Shor, the named inventors of the '189 patent. Based on this sworn testimony, it is my opinion that Yvan Leclerc and Stephen Lau conceived of and reduced to practice the TerraVision system well before the earliest corroborated conception date that Skyline could possibly claim for the '189 patent. [*See* Shor Depo. at 321:17-324:5; *see also* SKY 001597.4.] The TerraVision system was both known by and used by others in the United States before this date, and it was described and published in the MAGIC Final Report (Ex. E) the TerraVision Technical Note (Ex. F), and the MAGIC IEEE Article (Ex. G) before this date.

2. *The TerraVision Source Code*

32. In my opinion, the TerraVision source code describes the operation of the TerraVision system on or about April 1996.

33. Below, I briefly describe certain relevant functions of the TerraVision source code. For convenience, I have described each function according to the file in which it is located.

34. TerraVision.c [src\TerraVision\TerraVision] includes the following functions:

a. **TerraVisionSpawnProcs**: This function spawns various threads, including the tile visibility thread, the rendering thread, and the tile requesting thread.

b. **TerraVisionRender**: This function executes as the rendering thread.

35. TsTsm.c [src\TerraVision\libTileSvc] includes the following functions:

a. **TSMSpawnThreads**: This function spawns a tile receiver thread for each server.

36. Visible.c [src\TerraVision\TerraVision] includes the following functions:

a. **GenerateVisible**: This function executes as the tile visibility thread. It loops, iteratively calling ThreeDWidgetGenerateVisible to determine how to generate the next frame.

b. **GenerateRequests**: This function calls ThreeDWidgetGenerateRequests to generate a quadtree. GenerateRequests then creates a list of the tiles in that quadtree in sorted order, from coarsest to finest resolution.

c. **TsMakeRequest**: This function executes as the tile requesting thread. It determines the rate at which requests to download tiles should be generated and loops, iteratively calling GenerateAndSendRequests.

d. **GenerateAndSendRequests**: This function waits for an alarm based on the download request update rate and then calls GenerateRequests to create a sorted list of the desired tiles (tiles up to the appropriate resolution) from coarsest to finest resolution. It then loops through this list in order, either placing a tile on a download request list if it is not already resident in the cache (by calling tsmReqTile), or updating the usage time on the tile if it is already resident (by calling TileMgrUpdateTile). If any tiles were placed on the download request list, then tsmStopReqTiles is called.

e. **ParseQuadTree**: This function takes a quadtree and returns a list of the "leaf" tiles that are resident in memory. It is called from ThreeDWidgetGenerateVisible.

37. ThreeDWidget1.c [src\TerraVision\TerraVision] includes the following functions:

a. **ThreeDWidgetGenerateVisible**: This function calls ThreeDWidgetCalcVisibility to generate a quadtree of visible tiles at up to the appropriate resolution based on the user's current view matrix. It then calls ParseQuadTree to determine the "leaf" tiles in the quadtree that are resident in memory and calls ThreeDWidgetCreateRenderPrimitive to create the mesh that will be rendered during the next frame.

b. **ThreeDWidgetGenerateRequests**: This function also calls ThreeDWidgetCalcVisibility to generate a quadtree of requested tiles at up to the appropriate resolution, based on a "bloated" view matrix (i.e., a view matrix specifying an expanded version of the view frustum used by ThreeDWidgetGenerateVisible).

c. **ThreeDWidgetCalcVisibility**: This recursive function creates a quadtree whose nodes are contained within the frustum defined by the passed view matrix up to the appropriate resolution based on that view matrix.

d. **ThreeDWidgetCreateRenderPrimitive**: The function creates the underlying polygonal mesh that will be used to draw the scene, based on the view matrix and the visible tiles.

38. tsmTileIO.c [src\tsmApi] includes the following functions:

a. **tsmReqTile**: This function places on a download request list a record for a tile containing specified coordinates at a specified resolution.

b. **tsmStopReqTiles**: This function is called after all calls to tsmReqTile have been made for a frame. It calls a function specific to the kind of protocol that will be used to download the tiles, including tsmStopReqTiles_iss and tsmStopReqTiles_web.

c. **tsmGetNextTileHeader**: This function gets the next tile's header by calling a function specific to the kind of protocol being used to download the tile, including tsmGetNextTileHeader_iss and tsmGetNextTileHeader_web.

d. **tsmGetNextTileData**: This function gets the next tile's data by calling a function specific to the kind of protocol being used to download the tile, including tsmGetNextTileData_iss and tsmGetNextTileData_web.

39. tsmTileIO_iss.c [src\tsmApi] includes the following functions:

a. **tsmStopReqTiles_iss**: This function calls lower-level functions that send the tile requests to a remote server.

b. **tsmGetNextTileHeader_iss**: This function calls a lower-level function to read a tile's header from a remote server.

c. **tsmGetNextTileData_iss**: This function calls a lower-level function to read a tile's data from a remote server.

40. tsmTileIO_web.c [src\tsmApi] includes the following functions:

a. **tsmGetTile_web**: This function calls tsmHttpUrlToBuffer to get the tile through HTTP.

b. **tsmStopReqTiles_web**: This function does very little.

c. **tsmGetNextTileHeader_web**: This function does very little.

d. **tsmGetNextTileData_web**: This function calls tsmGetTile_web.

41. tsmUrlToFile.c [src\tsmApi] includes the following functions:

a. **tsmHttpUrlToBuffer**: This function reads all data at a remote server's URL into a passed buffer.

42. TsTileStruct.c [src\TerraVision\libTileSvc] includes the following functions:

a. **TsImageServerReader**: This function executes as the tile receiver thread. It loops, reading tiles by calling TsReadTileFromTSM.

b. **TsReadTileFromTSM**: This function reads a tile by calling tsmGetNextTileHeader and tsmGetNextTileData.

3. *Comparing TerraVision to the Asserted Claims of the '189 Patent*

43. I have included an Exhibit V identifying relevant disclosures in (1) the TerraVision source code, (2) the MAGIC Final Report, (3) the TerraVision Technical Note and (4) the MAGIC IEEE Article. Based on these disclosures, and other secondary materials and testimony describing the TerraVision system referenced herein, it is my opinion that the TerraVision system anticipates and/or renders obvious every limitation of claims 1-3, 5, 7-9, 11-14, 16, 18-19, and 21-24 of the '189 patent. In addition, it is my opinion that each of the identified references separately and independently anticipates and/or renders obvious every limitation of claims 1-3, 5, 7-9, 11-14, 16, 18-19, and 21-24 of the '189 patent.

a. *Claim 1, Preamble:* **A method of providing data blocks describing three-dimensional terrain data to a renderer, the data blocks belonging to a hierarchical structure which includes blocks at a plurality of different resolution levels**

44. The preamble of claim 1 recites a "method of providing data blocks describing three-dimensional terrain data." The Court has construed "data block describing three-dimensional terrain" to mean "a block or collection of data or digital information that represents or describes a section of three-dimensional terrain at a particular resolution level and that

includes any additional data overlaid on the digital image of the terrain, such as altitude, labels or optional objects." [Ex. D-1 at 9-12.]

45. The TerraVision system discloses a "method of providing data blocks describing three-dimensional terrain." [*See* Ex. V.] TerraVision is a high-speed graphics application that allows a user to interact in real time with a synthetic 3D photo-realistic view of a large terrain. This application accesses "tiles" representing both elevation and image data. Tiles with elevation data are referred to as digital elevation model ("DEM") tiles and tiles with image data are referred to as orthographic image ("OI") tiles. Both DEM tiles and OI tiles are data blocks describing three-dimensional terrain data as construed by the Court. They are data blocks or collections of digital information that represent or describe a section of three-dimensional terrain, that is the "surface features of an area of land, an object, or a material, including color, elevation, and existing objects or structures on the land, object or material." Moreover, both DEM tiles and OI tiles are provided, as rendering of the terrain on the screen is accomplished by combining the DEM and OI tiles for the selected area at the appropriate resolution, as illustrated in Figure 4 of the MAGIC IEEE Article:

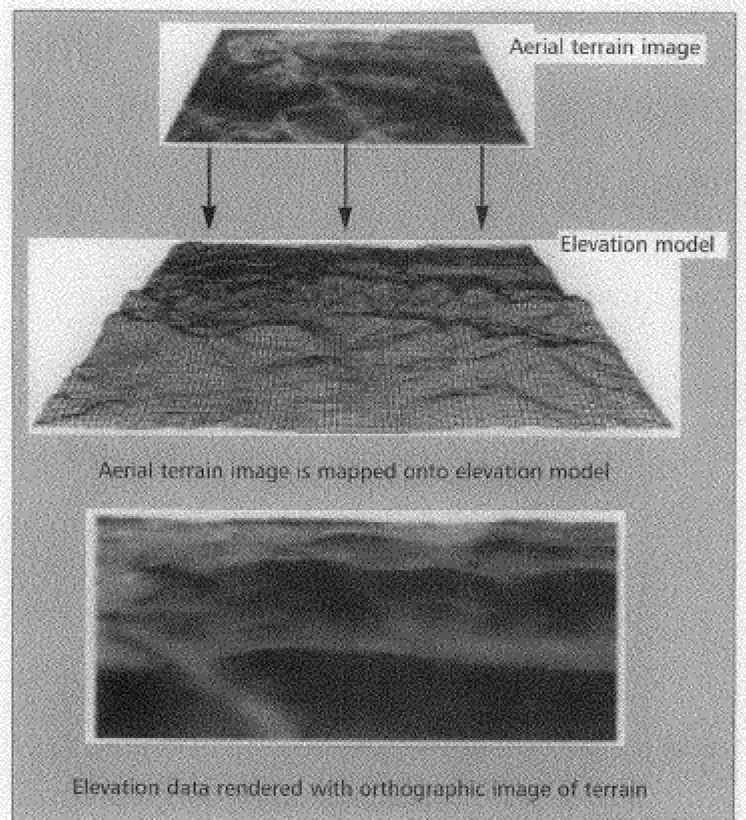

Figure 4. *Mapping an ortho-image onto its digital elevation model. (Source: SRI International)*

[Ex. G. at GOOG 350.] Further, additional objects can be overlaid on the terrain, including buildings and vehicles.

46. The preamble of claim 1 further indicates that the "data blocks describing three-dimensional terrain" are provided to a "renderer." The Court has construed the "renderer" to be

a "software and/or hardware object that performs at least the following functions: (1) determining and providing to another object the required coordinates in the terrain along with a respective resolution level; (2) receiving the data blocks corresponding to the specified coordinates; and (3) using the received data blocks to display a three-dimensional image." [Ex. D-1 at 26-32.] Based on the testimony of Stephen Lau, co-inventor of TerraVision, TerraVision had a renderer as construed by the Court:

> Q. And the next term is "renderer," which the Court has construed as "software and/or hardware object that performs at least the following functions: One, determining and providing to another object the required coordinates in the terrain along with a respective resolution level; two, receiving the data blocks corresponding to the specified coordinates; and three, using the received data blocks to display a three-dimensional image." Was that found in the TerraVision system?
>
> A. Yes.
>
> Q. And what was it called in the TerraVision system?
>
> A. Renderer. The renderer.

[Lau Depo. at 190:6-20.]

47. Further, in my opinion, the TerraVision system discloses a "renderer" as construed by the Court. [*See* Ex. V.] The TerraVision source code confirms that the TerraVision system has software that performs the functions of (1) determining and providing to another object the required coordinates in the terrain along with a respective resolution level; (2) receiving the data blocks corresponding to the specified coordinates; and (3) using the received data blocks to display a three-dimensional image. Specifically, ThreeDWidgetGenerateVisible calls ThreeDWidgetCalcVisibility to create a quadtree identifying all the visible tiles within the frustum up to a resolution appropriate for the view matrix. Each tile is identified by an *x* coordinate, a *y* coordinate and a resolution level. ParseQuadTree (another object) receives the quadtree from ThreeDWidgetGenerateVisible (part of the renderer). ParseQuadTree then provides a list of the "leaf" tiles resident in memory to ThreeDWidgetGenerateVisible, which uses these tiles to create the underlying polygonal mesh that will be used to display the scene. Finally, the rendering thread will actually display the three-dimensional image based on the tiles provided.

48. Moreover, to the extent that Skyline seeks to construe the term "renderer" more broadly than the Court's construction in an attempt to capture Google Earth, then in my opinion, the MAGIC Final Report, the TerraVision Technical Note, and the MAGIC IEEE Article also disclose a "renderer." [*See* Ex. V.] In particular, these references all disclose a search algorithm used to identify all the visible tiles at the appropriate resolution level. They further disclose that a small fraction of the available tiles are stored in local cache so that TerraVision is able to display a new view at any time, no matter how quickly the user moves. Based on these disclosures, a person of ordinary skill in the art would understand that the renderer must receive these data blocks from local memory and use them to display a three-dimensional image.

49. The preamble of claim 1 further recites "data blocks belonging to a hierarchical structure which includes blocks at a plurality of different resolution levels." The Court has construed "data blocks belonging to a hierarchical structure" as "data blocks that are organized into multiple levels of resolution, whereby each level contains data blocks at the same resolution, and each successive level contains data blocks of a higher resolution than those in the preceding level."

50. The TerraVision system discloses "data blocks belonging to a hierarchical structure which includes blocks at a plurality of different resolution levels," as that phrase has been construed by the Court. [*See* Ex. V.] TerraVision uses a multiresolution hierarchy or pyramid of increasingly lower-resolution representations of the equal-sized tiles. Each level in the pyramid is at half the resolution of the previous level. This is illustrated in Figure 3 of the MAGIC IEEE Article:

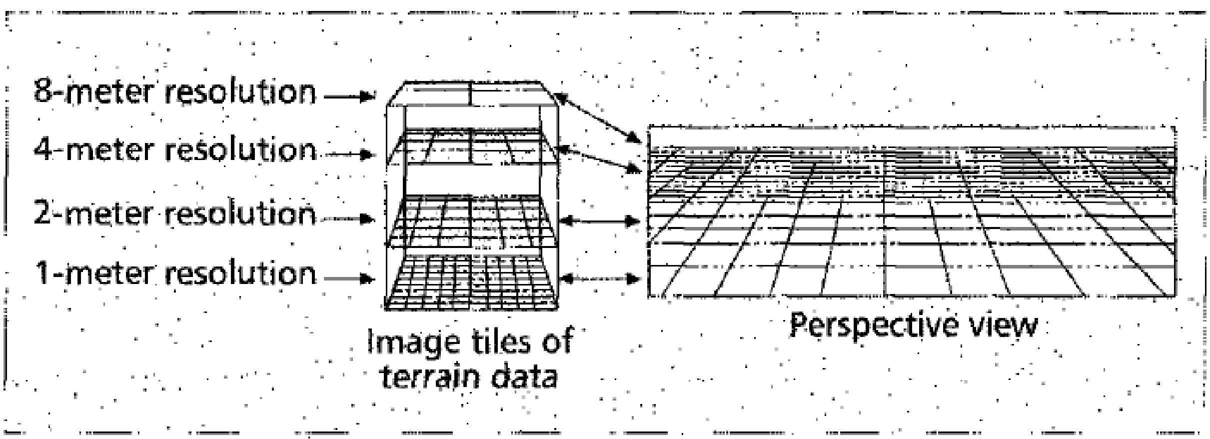

■ Figure 3. *Relationship between tile resolutions and perspective view. (Source: SRI International)*

[Ex. G at GOOG 349.]

51. The TerraVision source code also confirms that TerraVision had "data blocks belonging to a hierarchical structure which includes blocks at a plurality of different resolution levels," as that phrase has been construed by the Court. The ThreeDWidgetCalcVisibility function in the TerraVision source code creates a quadtree identifying tiles at a plurality of different resolution levels. It can do this because the tiles on the remote servers have been organized as pyramids. Thus, the quadtree created by ThreeDWidgetCalcVisibility and the pyramids on the remote servers are organized into multiple levels of resolution, where each level has tiles at the same resolution, and each successive level has tiles of a higher resolution than those in the preceding level.

b. *Claim 1:* receiving from the renderer one or more coordinates in the terrain along with indication of a respective resolution level

52. The Court has construed "coordinates in the terrain" as "a set of numerical values that identifies a particular location in the terrain," and "terrain" as "the surface features of an area

of land, an object, or a material, including color, elevation, and existing objects or structures on the land, object or material." [Ex. D-1 at 17-23.] The Court has also clarified that "receiving from the renderer" means "something distinct from the renderer receiving from the renderer." [Ex. D-2 at 8-10.]

53. The TerraVision system discloses the step of "receiving from the renderer one or more coordinates in the terrain along with indication of a respective resolution level." [*See* Ex. V.] As discussed above, ThreeDWidgetGenerateVisible calls ThreeDWidgetCalcVisibility to create a quadtree with all of the visible tiles within the frustum up to a resolution appropriate for the view matrix. Each tile is identified by an *x* coordinate, a *y* coordinate and a resolution level. ParseQuadTree (another object) receives this quadtree from ThreeDWidgetGenerateVisible.

54. This is confirmed by the MAGIC Final Report, the TerraVision Technical Note, and the MAGIC IEEE Article. Tiles are identified by an *x* coordinate, a *y* coordinate and a resolution level. The (*x*,*y*) coordinates are a set of numerical values that identify the location in the terrain. Thus, TerraVision references tiles using "coordinates in the terrain," as defined by the Court. A high-speed search algorithm is used to find all the visible tiles at the appropriate resolution level.

c. *Claim 1:* providing the renderer with a first data block which includes data corresponding to the one or more coordinates, from a local memory

55. The Court has construed "first data block" as "the first data block provided to the renderer from the local memory corresponding to the specified coordinates." [Ex. D-1 at 15-17.] The Court has construed "local memory" as "memory easily accessible to the user's processor, either because it is physically part of the processor or is attached directly thereto, and distinct from the memory of the remote server from which data must be downloaded." [*Id.* at 32-34.] The Court has also clarified that "providing the renderer" means "something distinct from the renderer providing to the renderer." [Ex. D-2 at 8-10.]

56. The TerraVision system discloses "providing the renderer with a first data block which includes data corresponding to the one or more coordinates, from a local memory." [*See* Ex. V.] ThreeDWidgetGenerateVisible calls ParseQuadTree, which identifies a "first data block which includes data corresponding to the one or more coordinates, from a local memory." Specifically, ParseQuadTree (something distinct from the renderer) identifies the visible tiles resident in memory and provides them to the renderer.

57. This is confirmed by the MAGIC Final Report, the TerraVision Technical Note, and the MAGIC IEEE Article. TerraVision keeps a small fraction of the available tiles in a local cache (i.e., in local memory). Moreover, by using local memory (i.e., by providing the first data block to the renderer), TerraVision is able to display a new view at any time, no matter how quickly the user moves, although that view may not be at the desired resolution. Both the visible tiles in memory and the tiles to be requested are identified using a "coarse-to-fine" strategy. This strategy manifests itself by the affected portion of the rendered image appearing "fuzzy" for a brief period of time before additional data blocks are downloaded. [*See also* Ex. H (video illustrating "coarse-to-fine" strategy).]

d. ***Claim 1:*** **downloading from a remote server one or more additional data blocks at a resolution level higher than the resolution level of the first data block which include data corresponding to the one or more coordinates if the provided block from the local memory is not at the indicated resolution level**

58. The Court has construed "downloading" to mean "requesting over a network from a separate computer and receiving on a local computer." [Ex. D-2 at 4-8.] The Court has construed "downloading … if the provided data block from the local memory is not at the indicated resolution level" to mean "downloading … upon some determination that the block provided from local memory is not at the indicated resolution level." [Ex. D-2 at 10-12.]

59. The TerraVision system discloses "downloading from a remote server one or more additional data blocks at a resolution level higher than the resolution level of the first data block which include data corresponding to the one or more coordinates." [*See* Ex. V.] The tile requesting thread in TerraVision requests tiles from a "remote server." The tile receiver thread receives tiles from the remote server. Moreover, the coarsest tiles are placed in the request list first. This placement is important because it means that, on average, the coarsest tiles (covering the largest spatial extent) will be received first. Indeed, in some cases, the coarsest tiles will always be received first. Thus, requests for tiles are ordered by resolution level and progressively higher resolution level data blocks are downloaded as needed to render the view. This is also described as a "coarse-to-fine" strategy. This strategy is illustrated in the TerraVision videos. [Ex. H.] When TerraVision first displays a new area, the image seems "out of focus" because the high-resolution tiles needed are not available in local memory and TerraVision is forced to use lower resolution tiles. [*Id.* at 173:9-11.] This "fuzzy" image becomes "progressively better focused" as TerraVision requests the higher-resolution tiles from the server and renders them on the display. [*Id.* at 173:9-21.]

60. Moreover, the TerraVision system discloses "downloading … if the provided block from the local memory is not at the indicated resolution level," as construed by the Court. ThreeDWidgetGenerateRequests calls ThreeDWidgetCalcVisibility to generate a quadtree of requested tiles at up to the appropriate resolution based on a "bloated" view matrix (i.e., a view matrix specifying an expanded version of the view frustum used by ThreeDWidgetGenerateVisible). This quadtree is passed to GenerateAndSendRequests as a sorted list, in coarse-to-fine resolution order. The tiles in the list include those in local memory, as well as those on a remote server. GenerateAndSendRequests then loops through the list, in coarse-to-fine order, to determine whether or not each tile is in local memory. If a tile in the list is already in local memory, then GenerateAndSendRequests will cause the usage time on the tile to be updated, so as to prevent the tile from being deleted from local memory, and it will not attempt to download that tile from a remote server. When GenerateAndSendRequests encounters a tile that is not in local memory (i.e., as it continues its coarse-to-fine traversal), then it will place that tile on a download request list. The download request list is processed by the tile set manager, which actually downloads the tiles from a remote server. The tile set manager supports several alternative methods for downloading tiles, including the use of ISS protocols and the use of HTTP to access tiles addressed as URLs. GenerateAndSendRequests will request that additional tiles at a resolution level higher than the resolution level of the first tile provided

from local memory be downloaded *if* the provided tile from local memory is not at the indicated resolution level. TerraVision thus downloads tiles from the remote server "upon some determination that the block provided from local memory is not at the indicated resolution level."

61. This is confirmed by the MAGIC Final Report and the TerraVision Technical Note. In particular, the MAGIC Final Report states that "[t]here is a 'tile prediction' thread for predicting the user's movement and determining the set of tiles that will be visible in the predicted viewpoint, and a 'tile requesting' thread for requesting those tiles from the ISS (eliminating those that are already in memory)." [Ex. E at GOOG 364.] This step of "eliminating those that are already in memory" is a reference to the determination made by TerraVision of whether the block provided from local memory is at the indicated resolution level—TerraVision will only download if the block in memory is not at the indicated resolution level. Furthermore, the TerraVision Technical Note provides that "[t]he truncated quad tree is then traversed in a breadth-first fashion. For each node encountered, a test is made to see if the corresponding color tile is currently in memory (since the first few top levels are always in memory, the test is not necessary at these levels). If the tile is not currently in memory, the coordinates of the tile are placed at the bottom of the list of tiles to pre-fetch." [Ex. F at GOOG 388.] Again, this is a reference to the same determination of whether the block provided from local memory is at the indicated resolution level. Thus, both the MAGIC Final Report and the TerraVision Technical Note specifically disclose the "downloading … if" limitation as construed by the Court.

62. Moreover, the MAGIC IEEE Article in combination with any of the other TerraVision references would at least render obvious the "downloading … if" limitation. While the MAGIC IEEE Article does not specifically disclose a "determination" of whether the block provided from local memory is at the indicated resolution level, it does describe the same system as the TerraVision source code, the MAGIC Final Report and the TerraVision Technical Note, and it would have been obvious to combine these references. Indeed, the TerraVision Technical Note is explicitly referenced in the MAGIC Final Report and the MAGIC IEEE Article, and a person of ordinary skill in the art would have been motivated to combine these references.

e. *Claim 2:* A method according to claim 1, wherein downloading the one or more additional data blocks comprises downloading the blocks from a succession of resolution levels, from the level immediately higher than the resolution level of the first block up to the maximal existent resolution level on the server not above the indicated resolution level.

63. Claim 2 includes all the limitations of claim 1, and a further limitation relating to "downloading the blocks from a succession of resolution levels." The Court has construed "succession of resolution levels" to mean "in order of increasing resolution level." [Ex. D-2 at 12-13.]

64. As I understand the Court's construction, the additional data blocks must both be requested in order of increasing resolution level and received in order of increasing resolution level. The TerraVision system requests blocks in a coarse-to-fine order (i.e., from the level immediately higher than the resolution level of the first block up to the maximal existent

resolution level on the server not above the indicated resolution level). In particular, GenerateAndSendRequests makes calls to tsmReqTile in coarse-to-fine order to generate a tile download request list that is sorted in order of increasing resolution level. The TerraVision system also receives blocks in a coarse-to-fine order at least when it uses the HTTP protocol to access tiles addressed as URLs. In TerraVision, this interface is implemented synchronously, meaning that TerraVision will wait to receive a requested block before requesting the next block. This is shown, for example, by tsmGetTile_web and tsmHttpUrlToBuffer. [*See also* Lau Depo. Exs. 99, 200 (addressing development of TSM API, including use of the HTTP protocol to access tiles over the World Wide Web).] In addition, even when TerraVision uses the ISS protocol to access tiles, on average, the coarsest tiles (covering the largest spatial extent) will be received first. This ISS interface, however, will not wait to receive a requested block before requesting the next data block. Accordingly, based on my understanding of the Court's construction, TerraVision discloses this limitation of claim 2 at least when it uses the HTTP protocol. [*See* Ex. V.] Moreover, TerraVision at least renders obvious this limitation when it uses the ISS protocol. [*See id*.] TerraVision already suggests requesting data blocks in order of increasing resolution level, identifies the advantages of receiving blocks in the same order and implements alternative protocols which ensure receipt of blocks in order of increasing resolution level. Moreover, to the extent that Skyline seeks to construe claim 2 more broadly in an attempt to capture Google Earth, it is my opinion that under any broader construction, TerraVision would disclose this limitation of claim 2 when it uses either the HTTP protocol or the ISS protocol.

f. *Claim 3*

65. Claim 3 is an independent claim. However, the limitations of claim 3 are substantially identical to limitations of claim 1. I address the differences below.

66. Claim 3 includes further language in the "receiving" step referring to the "plurality of coordinates being included in a plurality of respective distinct blocks." The Court has construed this language to mean "several coordinates, where each set of coordinates is contained within one block in a set composed of data blocks that are distinct from one another." [Ex. D-2 at 14-16.]

67. The Court's construction appears to encompass two different situations. First, where the renderer calls for a set of blocks using the (*x*,*y*) coordinates and resolution level of each block. And second, where the renderer calls for a set of blocks using coordinates of the boundaries of the necessary areas and a resolution level. *See* '189 patent, col. 14:10-16. Thus the coordinates may refer to one block or they may refer to many blocks. [Ex. D-2 at 16 ("I construe this language more broadly, allowing a one-to-many relationship between the blocks and the coordinates").] As discussed above with respect to claim 1, the renderer in TerraVision calls for a set of blocks using the (*x*,*y*) coordinates and resolution level of each block, and thus falls within the Court's construction of claim 3. [*See* Ex. V.]

68. There are some minor variations in the language of the "providing" step of claim 3. However, these differences do not change my analysis: the TerraVision system practices this step for the same reasons discussed above with respect to claim 1.

69. Claim 3 also includes further language in the "downloading" step that "blocks of lower resolution levels are downloaded before blocks of higher resolution levels." The TerraVision system requests blocks in a coarse-to-fine order (i.e., blocks of lower resolution levels are downloaded before blocks of higher resolution levels). In particular, GenerateAndSendRequests makes calls to tsmReqTile in coarse-to-fine order to generate a tile download request list that is sorted in order of increasing resolution level. The TerraVision system also receives blocks in a coarse-to-fine order at least when it uses the HTTP protocol to access tiles addressed as URLs. In TerraVision, this interface is implemented synchronously, meaning that TerraVision will wait to receive a requested block before requesting the next block. This is shown, for example, by tsmGetTile_web and tsmHttpUrlToBuffer. [*See also* Lau Depo. Exs. 99, 200 (addressing development of TSM API, including use of the HTTP protocol to access tiles over the World Wide Web).] In addition, even when TerraVision uses the ISS protocol to access tiles, on average, the coarsest tiles (covering the largest spatial extent) will be received first. This ISS interface, however, will not wait to receive a requested block before requesting the next data block. Accordingly, based on my understanding of the Court's construction, TerraVision discloses this limitation of claim 3 at least when it uses the HTTP protocol. [*See* Ex. V.] Moreover, TerraVision at least renders obvious this limitation when it uses the ISS protocol. [*See id.*] TerraVision already suggests requesting data blocks in order of increasing resolution level, identifies the advantages of receiving blocks in the same order and implements alternative protocols which ensure receipt of blocks in order of increasing resolution level. Moreover, to the extent that Skyline seeks to construe claim 3 more broadly in an attempt to capture Google Earth, it is my opinion that under any broader construction, TerraVision would disclose this limitation of claim 3 when it uses either the HTTP protocol or the ISS protocol.

g. *Claim 5*

70. Claim 5 is an independent claim. However, the limitations of claim 5 are substantially identical to limitations of claim 1 and 3. I address the differences below.

71. Claim 5 includes further language in the "downloading" step that "the blocks are downloaded according to the order in which the coordinates were provided." The TerraVision system requests blocks in a coarse-to-fine order (i.e., blocks are downloaded according to the order in which the coordinates were provided). In particular, GenerateAndSendRequests makes calls to tsmReqTile in coarse-to-fine order to generate a tile download request list that is sorted in order of increasing resolution level. The TerraVision system also receives blocks in a coarse-to-fine order at least when it uses the HTTP protocol to access tiles addressed as URLs. In TerraVision, this interface is implemented synchronously, meaning that TerraVision will wait to receive a requested block before requesting the next block. This is shown, for example, by tsmGetTile_web and tsmHttpUrlToBuffer. [*See also* Lau Depo. Exs. 99, 200 (addressing development of TSM API, including use of the HTTP protocol to access tiles over the World Wide Web).] In addition, even when TerraVision uses the ISS protocol to access tiles, on average, the coarsest tiles (covering the largest spatial extent) will be received first. This ISS interface, however, will not wait to receive a requested block before requesting the next data block. Accordingly, based on my understanding of the Court's construction, TerraVision discloses this limitation of claim 5 at least when it uses the HTTP protocol. [*See* Ex. V.] Moreover, TerraVision at least renders obvious this limitation when it uses the ISS protocol. [*See id.*] TerraVision already suggests requesting data blocks in order of increasing resolution

level, identifies the advantages of receiving blocks in the same order, and implements alternative protocols that ensure receipt of blocks in order of increasing resolution level. Moreover, to the extent that Skyline seeks to construe claim 5 more broadly in an attempt to capture Google Earth, it is my opinion that under any broader construction, TerraVision would disclose this limitation of claim 5 when it uses either the HTTP protocol or the ISS protocol.

h. *Claim 7*

72. Claim 7 is an independent claim. However, the limitations of claim 7 are substantially identical to limitations of claim 1. I address the differences below.

73. Claim 7 includes further language in the "downloading" step requiring "downloading from a remote server excess blocks not currently needed by the renderer to fill up the local memory when not downloading blocks required by the renderer." The Court has construed the phrase "when not downloading blocks required by the renderer" to mean "during periods of time when the local computer, or a connection thereof, is not downloading data blocks in response to coordinates received from the renderer." [Ex. D-2 at 17-20.]

74. In the TerraVision system, the tile requesting and tile receiver threads download the visible data blocks, as well as excess blocks not currently needed by the renderer to fill up the local memory. ThreeDWidgetGenerateRequests generates a quadtree of requested tiles at up to the appropriate resolution, based on a "bloated" view matrix (i.e., a view matrix specifying an expanded version of the view frustum used by the tile visibility thread). This quadtree will include all the visible tiles identified by ThreeDWidgetGenerateVisible (i.e., the tiles needed by the renderer) as well as excess tiles from the bloated view matrix. GenerateAndSendRequests places both the visible tiles and the excess tiles on a download request list if they are not already in memory. The tile requesting and tile receiver threads will actually download these tiles from the remote server. This is confirmed by the MAGIC Final Report and the MAGIC IEEE Article. [*See* Ex. E at GOOG 364; Ex. G at GOOG 351 (also addressing assignment of download priorities).] The TerraVision Technical Note, however, includes further disclosures. It provides that "[u]sing an expanded field of view of the predicted viewpoint is a relatively simple mechanism that has the advantage that exactly the same code used for traversing the terrain quad tree for the current view can be used to create a truncated quad tree for the future view." [Ex. F at GOOG 388.] It further describes a "secondary list of 'tiles to pre-fetch if there's time'…." [Ex. F at GOOG 388.] In my opinion, at least the TerraVision Technical Note discloses "downloading from a remote server excess blocks not currently needed by the renderer to fill up the local memory when not downloading blocks required by the renderer." [*See* Ex. V.] The excess tiles are the tiles placed on the secondary list of tiles to pre-fetch if there's time. These tiles are only downloaded during periods of time when the local computer, or a connection thereof, is not downloading data blocks required by the renderer (i.e., "if there's time"). Furthermore, the TerraVision system, as well as the MAGIC Final Report and the MAGIC IEEE Article, would at least render this limitation obvious. [*Id.*] These references disclose downloading excess blocks not currently needed by the renderer to fill up the local memory. TerraVision, however, does not first download the tiles needed by the renderer and then download the excess data blocks—these blocks are all placed on the same download request list in coarse-to-fine order. A person of ordinary skill in the art nonetheless would have been motivated to download the excess tiles only after downloading the tiles need by the renderer as

suggested by the TerraVision Technical Note. Moreover, to the extent that Skyline seeks to construe claim 7 more broadly in an attempt to capture Google Earth, it is my opinion that under any broader construction, the TerraVision source code, the MAGIC Final Report and the MAGIC IEEE Article would also disclose this limitation of claim 7.

i. *Claim 8:* A method according to claim 7, wherein downloading the data blocks comprised downloading the blocks via the Internet.

75. Claim 8 includes all the limitations of claim 7, and a further limitation relating to "downloading the blocks via the Internet." The Court has construed the "Internet" to mean "[t]he publicly accessible network capable of relaying information via Internet Protocol, either alone or in conjunction with one or more other protocols, but not including a wholly self-contained private network of devices communicating only with each other." [Ex. D-2 at 20-22.] The Court further noted that "this construction is not intended to exclude networks such as Inte[r]net2, which are built on separate physical infrastructure, but are essentially updated and experimental versions of the current Internet." [*Id.* at 22.]

76. The TerraVision system performs this further limitation of claim 8. [*See* Ex. V.] In particular, as disclosed in the MAGIC Final Report, the TerraVision system was demonstrated over the Internet at the Jet Propulsion Laboratory prior to May 1996. In addition, Stephen Lau testified that he routinely ran TerraVision over the Internet from SRI to remote servers at Lawrence Berkeley National Labs. [Lau Depo. at 38:10–22, 281:25–283:1.] Mr. Lau also testified that he demonstrated TerraVision at SIGGRAPH 95 by connecting through the Internet to remote servers at Sioux Falls, South Dakota (Lau Depo. at 90:7–18), and even ran TerraVision from his home and hotel rooms using a 9600 baud modem (Lau Depo. at 37:4–38:5). [*See also* Clinger Decl. & Exs; Marke Clinger, "GraphicsNet '95: Integrated voice, video, graphics and data network using asynchronous transfer mode (ATM)," *ACM SIGGRAPH Computer Graphics*, 30(1), pp. 10-18 (Feb. 1996).] In addition, the TerraVision source code includes functions for accessing tiles over the Internet by using the HTTP protocol, including, for example, tsmGetTile_web and tsmHttpUrlToBuffer.

77. In addition, in my opinion, the MAGIC network is akin to the Internet2 (i.e., it was an updated and experimental version of the then-current Internet). The "I" in MAGIC stands for "Internet," and the MAGIC network was specifically designed as a testbed network. Furthermore, the MAGIC network was an ATM network. At the relevant time, ATM networks were gaining increasing acceptance as a high-speed Internet backbone. [*See, e.g.*, Exs. S-U.] Accordingly, in my opinion, the MAGIC network would fall within the Court's construction of "Internet," and use of TerraVision over the MAGIC network would also satisfy this further limitation of claim 8.

78. Furthermore, even assuming that the MAGIC network is a private network, it was also part of the broader Internet, not a "self-contained group of computers or other devices communicating only with one another." [*See* Ex. D-2 at 21.] For example, Stephen Lau testified that SRI was not on the MAGIC gigabit testbed, and therefore had to access the MAGIC network over the Internet. [Lau Depo. at 38:10–22, 281:25–283:1.] This is confirmed in the MAGIC IEEE Article. [Ex. G at GOOG 356 ("Proper testing of TerraVision and the ISS required high-

speed interconnectivity. However, SRI and LBNL, the respective developers of these components, did not have such connectivity.").] Accordingly, in my opinion, use of the MAGIC network over the broader Internet would also fall within the Court's construction.

79. Finally, there was a trend in the art to adapt 3D terrain visualization applications for use on the Internet, including the World Wide Web. This is illustrated, for example, by Rabinovich: "The advent of the World-Wide-Web suggests the running of this type of [terrain visualization] application over the Internet, in a client/server scenario." [Ex. Q at GOOG 9828.] Likewise, Potmesil describes a "WWW-based system … which allows users to view, search and post geographically-indexed information of the Earth," including "a 3D flight-simulator browser capable of continuous flight around the Earth." [Ex. R at GOOG 9762.] He states that "Internet-based computers and communications can be very effective in enhancing our ability to visualize and to search 3D environments in the great outdoors where we move, work, play and learn." [*Id.*] One of ordinary skill in the art would have been motivated to combine the World-Wide-Web-based strategies of Rabinovich and Potmesil with ATM-based terrain visualization applications such as TerraVision and T_Vision. Potmesil specifically references ATM-based terrain visualization applications, including T_Vision, and is directed to the same subject matter. [*Id.* at GOOG 9764.] Accordingly, in my opinion, claim 8 would have at least been rendered obvious by the combination of one of TerraVision or T_Vision with one of Rabinovich or Potmesil.

j. *Claim 9:* A method according to claim 7, wherein the renderer renders a view from a current viewpoint, and wherein downloading the excess blocks comprises filling the local memory with substantially all of the blocks surrounding a point in the terrain seen from the current viewpoint within a predetermined distance range.

80. Claim 9 includes all the limitations of claim 7, and a further limitation relating to downloading "substantially all the blocks surrounding a point in the terrain seen from the current viewpoint within a predetermined range." The Court construed this phrase to mean "substantially all of the data blocks covering terrain within a uniform predetermined distance from a point in the terrain that is seen from the current viewpoint." [Ex. D-2 at 22-26.]

81. As discussed above with respect to claim 7, TerraVision downloads excess tiles based on a "bloated" view matrix (i.e., a view matrix specifying an expanded version of the view frustum used by the tile visibility thread). In many typical cases, use of this bloated view matrix will result in substantially all of the blocks covering terrain within a uniform predetermined distance from a point in the terrain that is seen from the current viewpoint being downloaded.

82. Furthermore, downloading blocks surrounding a point in the terrain seen from the current viewpoint within a predetermined distance range was known in the art. This is shown, for example, by Potmesil, who discloses the downloading of excess data blocks using several different caching strategies, one of which is downloading blocks surrounding a point in the terrain seen from the current viewpoint within a predetermined distance range was known in the art. [*See* Ex. R at GOOG 9767.] He states: "The caching algorithm uses the user's current position, velocity, and acceleration to estimate where the user is moving and allocates new tiles there." [*Id.*] One of the "several implemented caching strategies" is "obtain as many tiles as

close to the view as possible; this may be used in low speeds (hovering) when the direction of flight is uncertain." [*Id.*] Potmesil also discloses filling the local memory with excess tiles: "The more disc and memory space the host machine has available, the more tiles can be brought into the cache and remain there." [*Id.*] One of ordinary skill in the art would have been motivated to combine the caching strategies of Potmesil with TerraVison as Potmesil specifically references ATM-based terrain visualization applications, including T_Vision, and is directed to the same subject matter. [*Id.* at GOOG 9764.] Accordingly, in my opinion, claim 9 would have at least been rendered obvious by the combination of one of TerraVision or T_Vision with Potmesil.

k. *Claim 11:* A method according to claim 9, wherein filling the local memory comprises filling the memory with substantially all the blocks within the range from a lower resolution level before downloading blocks of higher resolution levels.

83. Claim 11 includes all the limitations of claims 7 and 9, and a further limitation relating to filling the local memory with "substantially all the blocks within the range from a lower resolution level before downloading block of higher resolution level."

84. The TerraVision system requests blocks in a coarse-to-fine order (i.e., blocks of lower resolution levels are downloaded before blocks of higher resolution levels). In particular, GenerateAndSendRequests makes calls to tsmReqTile in coarse-to-fine order to generate a tile download request list that is sorted in order of increasing resolution level. The TerraVision system also receives blocks in a coarse-to-fine order at least when it uses the HTTP protocol to access tiles addressed as URLs. In TerraVision, this interface is implemented synchronously, meaning that TerraVision will wait to receive a requested block before requesting the next block. This is shown, for example, by tsmGetTile_web and tsmHttpUrlToBuffer. [*See also* Lau Depo. Exs. 99, 200 (addressing development of TSM API, including use of the HTTP protocol to access tiles over the World Wide Web).] In addition, even when TerraVision uses the ISS protocol to access tiles, on average, the coarsest tiles (covering the largest spatial extent) will be received first. This ISS interface, however, will not wait to receive a requested block before requesting the next data block. Accordingly, based on my understanding of the Court's construction, TerraVision discloses this limitation of claim 11 at least when it uses the HTTP protocol. [*See* Ex. V.] Moreover, TerraVision at least renders obvious this limitation when it uses the ISS protocol. [*See id.*] TerraVision already suggests requesting data blocks in order of increasing resolution level, identifies the advantages of receiving blocks in the same order, and implements alternative protocols that ensure receipt of blocks in order of increasing resolution level. Moreover, to the extent that Skyline seeks to construe claim 11 more broadly in an attempt to capture Google Earth, it is my opinion that under any broader construction, TerraVision would disclose this limitation of claim 11 when it uses either the HTTP protocol or the ISS protocol.

k. *Claim 12:* An apparatus [of claim 1]

85. The limitations of claim 12 are substantially identical to those in claim 1, except that claim 12 is directed to an "apparatus" rather than a "method."

86. The TerraVision system is software designed to be run on a computer (i.e., an apparatus). [*See* Ex. E at GOOG 362 ("TerraVision was designed to run on certain Silicon Graphics (SGI) workstations…."); Ex. F at GOOG 385 ("Given a graphics station with an architecture similar to [that] of the popular Silicon-Graphics workstations…."); Ex. G at GOOG 350 ("High frame rates are achieved by using a local very-high-speed rendering engine, an SGI Onyx….").] In my opinion, there is nothing in claim 12 that would somehow read out an SGI workstation. Claim 12 is generally drawn to an "apparatus" with a "local memory," a "communications link" and a "processor"—elements found in the SGI workstations used in the TerraVision system.

87. Moreover, the early to mid 1990s SGI workstations used by SRI in the development of TerraVision and described in the MAGIC Final Report (Onyx with four 150 MHz R4400 CPUs), MAGIC Q4 1994 Quarterly Report (Indigo), and the Lau deposition (Indigo with R4000 or R4400 CPU, Indigo2) were comparable to or less powerful than many off-the-shelf PCs running Microsoft Windows available by February 1999, when the '189 patent was filed.

l. *Claim 13:* An apparatus [of claim 2]

88. The limitations of claim 13 are substantially identical to those of claims 1 and 2, except that claim 13 is directed to an "apparatus" rather than a "method." For the reasons stated with respect to claims 1, 2, and 12, in my opinion, TerraVision anticipates and/or renders obvious this claim.

m. *Claim 14:* An apparatus [of claim 3]

89. The limitations of claim 14 are substantially identical to those of claim 3, except that claim 14 is directed to an "apparatus" rather than a "method." For the reasons stated with respect to claims 3 and 12, in my opinion, TerraVision anticipates and/or renders obvious this claim.

n. *Claim 16:* An apparatus [of claim 5]

90. The limitations of claim 16 are substantially identical to those of claim 5, except that claim 16 is directed to an "apparatus" rather than a "method." For the reasons stated with respect to claims 5 and 12, in my opinion, TerraVision anticipates and/or renders obvious this claim.

o. *Claim 18:* An apparatus [of claim 7]

91. The limitations of claim 18 are substantially identical to those of claim 7, except that claim 18 is directed to an "apparatus" rather than a "method." For the reasons stated with respect to claims 7 and 12, in my opinion, TerraVision anticipates and/or renders obvious this claim.

p. *Claim 19:* An apparatus [of claim 9]

92. The limitations of claim 19 are substantially identical to those of claims 7 and 9, except that claim 19 is directed to an "apparatus" rather than a "method." For the reasons stated with respect to claims 7, 9, and 12, in my opinion, TerraVision anticipates and/or renders obvious this claim.

q. *Claim 21:* An apparatus [of claim 11]

93. The limitations of claim 21 are substantially identical to those of claims 7, 9, and 11, except that claim 21 is directed to an "apparatus" rather than a "method." For the reasons stated with respect to claims 7, 9, 11, and 12, in my opinion, TerraVision anticipates and/or renders obvious this claim.

r. *Claim 22:* An apparatus [of claim 8]

94. The limitations of claim 22 are substantially identical to those of claims 7 and 8, except that claim 22 is directed to an "apparatus" rather than a "method." Furthermore, the Court has construed this claim as requiring only a "connection to the Internet" as opposed to requiring that "some portion of the download [] happen over the broader Internet." [*See* Ex. D-2 at 22.] As discussed with respect to claim 8, not only was TerraVision demonstrated over the Internet, it was also connected to the Internet. For the reasons stated with respect to claims 7, 8, and 12, in my opinion, TerraVision anticipates and/or renders obvious this claim.

s. *Claim 23:* The method of claim 7, wherein the coordinates relate to the coordinates of a predetermined course of a flight vehicle.

95. Claim 23 includes all the limitations of claim 7, and a further limitation providing that "the coordinates relate to the coordinates of a predetermined course of a flight vehicle." The TerraVision system performs this further limitation of claim 7. [*See* Ex. V.] In particular, users of the TerraVision system could navigate through (i.e., "fly over") a representation of a 3D view of a large terrain. Furthermore, while the TerraVision system did not attempt to mimic the controls of an aircraft, it did allow the user to roam about the terrain at will, simulating the view one would experience as if flying in an aircraft. It also disclosed the possibility of imposing the flight limitations of a particular aircraft if one wanted to build a simulator based on TerraVision technology.

96. There was also a trend in the art to use terrain visualization systems for flight simulation. The Migdal patent is evidence of this trend: "As is well-known in graphics design, the display view can simulate flight by constantly moving the eyepoint along a terrain or landscape being viewed. Such flight can be performed automatically as part of a program application, or manually in response to user input such as a mouse or joystick movement." [Ex. N ('783 patent at col. 10:16-21).] A person of ordinary skill in the art would have been motivated to apply this well-known trend to other terrain visualization systems such as TerraVision and T_Vision. Accordingly, in my opinion, claim 23 would have at least been rendered obvious by the combination of one of TerraVision or T_Vision with the Migdal patent.

t. *Claim 24:* An apparatus [of claim 23]

97. The limitations of claim 24 are substantially identical to those of claims 7 and 23, except that claim 24 is directed to an "apparatus" rather than a "method." For the reasons stated with respect to claims 7, 12, and 23, in my opinion, TerraVision anticipates and/or renders obvious this claim.

3. *Response to the Manocha Report re: TerraVision*

98. I have also now had the opportunity to review the Rebuttal Expert Report of Dinesh Manocha, Ph.D., dated August 11, 2006. While I will not respond here to every point raised in this report, I did want to address certain statements by Dr. Manocha, particularly as they mischaracterize my opinions, the patent-in-suit and/or the prior art.

99. With respect to TerraVision, Dr. Manocha first suggests that the "objective and approach" of the TerraVision system was different than that of the system described in the '189 patent. However, this analysis ignores the claim language and attempts to import limitations into the claims that simply are not there.

100. Dr. Manocha, for example, repeatedly references use of the TerraVision system with the MAGIC network and asserts that this distinguishes the prior art from the asserted claims of the '189 patent. However, all but two of the claims require only a "communication link," and a link to the MAGIC network indisputably constitutes a "communication link." While claims 8 and 22 further require a connection to the Internet, this also does not exclude high speed networks such as the MAGIC network, "which are built on separate physical infrastructure, but are essentially updated and experimental versions of the current Internet." [Ex. D-2 at 22 n.17.] In addition, Dr. Manocha ignores evidence that the TerraVision system was used not just with the MAGIC network, but also with the broader Internet. [*See, e.g.*, Ex. E at GOOG 369.] The inventors of the TerraVision system also explicitly recognized that their "objective and approach" was applicable to slow networks, as well as high speed networks.

101. Dr. Manocha also repeatedly references use of the TerraVision system on a high-end workstation computer. However, again, there is nothing in the asserted claims that would exclude such a computer, which is an "apparatus" with a "local memory," a "communication link" and a "processor."

102. Finally, Dr. Manocha asserts that testimony from Michael Jones about alleged differences in the number of users of the Google system and the TerraVision system somehow distinguishes the TerraVision system from the '189 patent. However, again, there is nothing in the asserted claims that requires a certain number of users, and in fact, the claims are drawn to a method performed by a single user and a single apparatus.

103. The other principal distinction drawn by Dr. Manocha between TerraVision and the '189 patent relates to TerraVision's treatment of digital elevation model or "DEM" tiles. However, here too, Dr. Manocha is relying on an asserted distinction not found in the claims.

104. TerraVision uses both DEM tiles and orthographic image or "OI" tiles to render an image. There can accordingly be no real dispute that both DEM and OI tiles are "provid[ed] … to a renderer" in some manner, since they are in fact used by the renderer.

105. Moreover, both DEM tiles and OI tiles constitute "data blocks describing three-dimensional terrain," as construed by the Court. They are data blocks or collections of digital information that represent or describe a section of three-dimensional terrain, that is the "surface features of an area of land, an object, or a material, including color, elevation, and existing objects or structures on the land, object or material." If a data block (whether an OI tile or a DEM tile) is not in local memory, then the TerraVision system will download "one or more additional data blocks."

106. According to Dr. Manocha, the claims require more—they require that even if a DEM data block is already in local memory, then the system must download both the OI data block and the DEM data block together. That is not required. Only the data blocks not already in local memory need to be downloaded. It also makes no difference whether the downloaded blocks are OI tiles or DEM tiles, as both are data blocks describing three-dimensional terrain. Indeed, to the extent that Dr. Manocha is suggesting that there must be some data block with both image and elevation data, then he is taking a position inconsistent with his own infringement report—Google Earth, like TerraVision, treats image data blocks and elevation data blocks separately, and does not download them as a single package.

107. Moreover, as of at least April 1995, both OI tiles and DEM tiles were requested and downloaded from the remote server in TerraVision. [*See* Lau Depo. Ex. 99; Ex. F at GOOG 388.] DEM tiles were downloaded from the remote server when the user selected a data set, and then stored in local cache memory since there was only a relatively small number of DEM tiles. [*Id.*] In fact, DEM tiles were requested in coarse-to-fine order in TsRequestDems, and the download list was created using the exact same tsmReqTile function used for OI tiles. Furthermore, the inventors of TerraVision disclosed alternative systems whereby DEM tiles as well as OI tiles would be downloaded using their coarse-to-fine strategy. For example, in the MAGIC Final Report, the authors make no distinction between DEM and OI tiles, suggesting only that a "small fraction of the available tiles [be kept] in a local cache." [Ex. E at GOOG 364; *see also* Ex. G (MAGIC IEEE Article).] This disclosure (along with the other disclosures regarding TerraVision) was more than sufficient to enable a person of ordinary skill in the art to make the invention. More fully implementing the coarse-to-fine strategy for DEM tiles as well as OI tiles would have involved minor modifications of the TerraVision source code, which already includes some relevant functions and variables for this purpose.

108. Dr. Manocha's suggestion that the inventors of the TerraVision system could not solve this problem is also unsupported—the inventors clearly conceived and reduced to practice the solution as early as 1993, i.e., "[t]he approach is to retrieve elevation data and specially processed aerial photographs from a high-speed Image Server System (ISS) in real time as they are needed for image generation." [Lau. Depo. Ex. 90 (January 26, 1993 Quarterly Report).] That they did not immediately implement this approach for DEM tiles in source code does not mean it could not be done by a person of ordinary skill in the art; only that it was not a priority for them.

109. Moreover, even assuming that Dr. Manocha is correct and that TerraVision's treatment of DEM tiles somehow takes TerraVision out of the literal scope of the claims, there is an explicit teaching in the TerraVision system to download DEM tiles as well as OI tiles "in real time as they are needed for image generation." Accordingly, at minimum, the TerraVision system would have rendered the claims obvious. I also disagree with Dr. Manocha's analysis of the secondary considerations.

110. Finally, Dr. Manocha misstates my opinions in certain respects. In particular, his contention that I have admitted that certain elements are missing in TerraVision is not correct. [*See* Manocha ¶¶ 40, 44, 48.] Further, his contention that my opinions are inconsistent in some manner between infringement and invalidity is also not correct. [*See* Mancha ¶¶ 36-39, 41-43.] Indeed, Dr. Manocha appears to ignore differences between Google Earth and the prior art in some cases. Finally, Dr. Manocha's criticism that I have not put my report in the form of a chart is also unwarranted—I analyzed and identified relevant disclosures in each of the TerraVision references on an element-by-element basis. However, for convenience, I have attached as Exhibit V to this report a chart setting out these disclosures separately.

4. *Conclusions*

111. Based on my comparison of asserted claims to the TerraVision system and to the individual TerraVision references, it is my opinion that the TerraVision system and references anticipate or render obvious all asserted claims of the '189 patent. Specifically, in my opinion, claims 1-3, 5, 7-9, 11-14, 16, 18-19, and 21-24 are anticipated and/or rendered obvious by the TerraVision system, the TerraVision source code, the MAGIC Final Report, the TerraVision Technical Note and the MAGIC IEEE Article.

112. Furthermore, a person of ordinary skill in the art would have been motivated to combine any of the TerraVision references, since they all refer to exactly the same project and system. Both the MAGIC Final Report and the MAGIC IEEE Article specifically reference the TerraVision Technical Note, and the MAGIC Final Report and the TerraVision Technical Note were posted on the same website (which also included further descriptions of the TerraVision system), providing additional motivation to combine these references. In addition, any alleged differences between asserted claims and the prior art are such that the claims would have been obvious at the time the invention was made to a person having ordinary skill in the art. In particular, the TerraVision achieved all of the objectives of the claimed inventions and any differences would have merely related to implementation details.

D. <u>The T_Vision System</u>

113. In my opinion, the T_Vision system itself, as well as certain documents describing this system, anticipates and/or renders obvious claims 1-3, 5, 7-9, 11-14, 16, 18-19, and 21-24 of the '189 patent. These documents include the Mayer patent and the Mayer German patent application.

1. *Overview of the T_Vision System*

114. T_Vision was a method and system for pictorial representation of space-related data, particularly geographic data of flat or physical objects. It could be and was used for visualizing topographic data.

115. T_Vision was a system that included:

- a "virtual globe… modeled from high resolution spatial data and textured with high resolution satellite images." [Ex. M at p. 1, ¶ 1.]

- a "database, a real time rendering system," and a "specific concept of seamless links between different levels of detail," allowing "continuous zooming from a global view down to recognizable features of only a few centimeters in size." [*Id.* at p. 1, ¶¶1-2.]

- a renderer that "calculates the currently needed data and predicts the needed data for future flightpath. Then it [the renderer] requests the data for a special location with an appropriate resolution." [Ex. L at RENDERER.htm, p. 1 of 2.]

- As a result: "The user thereby has full control over which information to view, at what time, and at which location." [Ex. M at p. 1, ¶ 4.]

116. It is my understanding, based on the documents and testimony I have reviewed regarding T_Vision, that ART+COM, and specifically Pavel Mayer, Axel Schmidt, Joachim Sauter and Gerd Grüneis, developed the T_Vision system in the mid 1990s. I understand that ART+COM is a company based in Berlin, Germany, that focuses on interactive digital media projects.

117. The T_Vision system was conceived and reduced to practice earlier than August 1995. This is shown by the fact that the T_Vision system was publicly demonstrated as a working system at least as early as October 1994 at the International Telecommunications Union meeting in Kyoto, Japan. [Ex. L at ITU.HTM, p. 1 of 1.] The T_Vision system was also publicly demonstrated at SIGGRAPH '95, which was held in Los Angeles, California, August 6–11, 1995. [Exs. K & L; *see also* Lau Depo. at 201:3-205:12.] I attended the SIGGRAPH '95 conference in Los Angeles, California, and recall seeing T_Vision being demonstrated in the Interactive Communities venue. The T_Vision system is also described in the SIGGRAPH '95 Multimedia CD-ROM / SIGGRAPH '95 Proceedings CD-ROM excerpts, which was available at the SIGGRAPH '95 Conference. [Ex. L.] The T_Vision system was not abandoned, suppressed or concealed; indeed, the T_Vision system was publicly demonstrated and patented, and information about this system was widely disseminated as detailed herein.

118. The T_Vision system was in public use and publicly known and used by ART+COM and others in the United States no later than August 1995. This is confirmed by the evidence above. [Exs. K & L; *see also* Lau Depo. at 201:3-205:12.] As stated above, I attended the SIGGRAPH '95 conference in Los Angeles, California, and recall seeing T_Vision being demonstrated in the Interactive Communities venue.

119. The Mayer German patent application is itself a printed publication, published no later than July 1997, more than one year before the filing date of the '189 patent. The application for the Mayer patent was filed in the United States on December 17, 1996, and issued on August 8, 2000. This December 1996 filing date is before the earliest priority date that Skyline has claimed or can claim based on the filing date of the application relating to the '189 patent or its alleged invention date.

120. Accordingly, it is my understanding that the T_Vision system, the Mayer German patent application and the Mayer patent constitute prior art with respect to the '189 patent.

2. *Comparing T_Vision to the Asserted Claims of the '189 Patent*

121. I have included an Exhibit W identifying relevant disclosures in (1) the SIGGRAPH '95 T_Vision Project materials, (2) the Mayer patent, and (2) the Mayer German patent application. Based on these disclosures, and other secondary materials and testimony describing the T_Vision system referenced herein, it is my opinion that the T_Vision system anticipates and/or renders obvious every limitation of claims 1-3, 5, 7-9, 11-14, 16, 18-19, and 21-24 of the '189 patent. In addition, it is my opinion that each of the identified references separately and independently anticipates and/or renders obvious every limitation of claims 1-3, 5, 7-9, 11-14, 16, 18-19, and 21-24 of the '189 patent

122. In the paragraphs below, I have not repeated the Court's construction of relevant claim terms or opinions relating to the interpretation of the '189 patent, which are set forth in detail above with respect to TerraVision. However, I incorporate these discussions herein by reference as they are equally applicable.

a. *Claim 1, Preamble:* **A method of providing data blocks describing three-dimensional terrain data to a renderer, the data blocks belonging to a hierarchical structure which includes blocks at a plurality of different resolution levels**

123. The T_Vision system practices and discloses a "method of providing data blocks describing three-dimensional terrain to a renderer." [*See* Ex. W.] This is illustrated, for example, in the T_Vision Project materials published at SIGGRAPH '95:

[Ex. L at TVISION.HTL, p. 1 of 3.] These materials disclose that the T_Vision system demonstrated at SIGGRAPH '95 provided a "virtual globe" that was "modeled from high resolution spatial data and textured with high resolution satellite images." [*Id; see also id.* at TERRABAS.HTM, p. 2 of 3 ("Currently the source data consists of around 10 GB of image and DEM data, covering the whole world in 4 km/pixel, USA and Europe in 1km, Japan in 50m, some areas in USA and Germany in 50m, and parts of Berlin and Tokyo down to 30cm.").] The Mayer patent further discloses, "[a] method and device for the pictorial representation of space-related data, for example, geographical data of the earth." [Ex. I ('897 patent at Abstract).] Correspondingly, the Mayer German patent application discloses: "a method and device for the pictorial representation of space-related data, especially geographical data of physical objects, for example, the earth." [Ex. J at GOOG 21372.] The "representation" referred to by the Mayer patent and by the Mayer German patent application may be "based on a three-dimensional geometrical model." [Ex. I ('897 patent at col. 8:56-57); Ex. J at GOOG 21377.] The T_Vision system used a database with "pairs of index and data files containing 128x128 pixel texture images (surface, clouds) and 16x16 point elevation data" (i.e., "data blocks describing three-dimensional terrain"). [Ex. L at TERRABAS.HTM, p. 1 of 3.] Likewise, the Mayer patent and the Mayer German patent application both refer to the subdivision of terrain data into "sections" representing both elevation and image data. [Ex. I ('897 patent at col. 8:28-42 (object divided into sections and each section of object has fixed address)); Ex. J at GOOG 21377.] The T_Vision system provides the data blocks to a renderer, which allows the user to visualize the terrain. [Ex. L. at RENDERER.HTM; *see also* Ex. I ('897 patent at Abstract ("For a screen representation of a view of the object according to a field of view of a virtual observer, the required data are called up and shown only in the resolution required for each individual section of the image.")); Ex. J at GOOG 21372 ("For a screen representation of a view of the object according to a field of view of a virtual observer, the required data are retrieved and represented only in the resolution required for every individual section of the image.").]

124. The T_Vision system discloses a "renderer" (at least as interpreted by Skyline), that is, software and/or hardware that performs the following functions: (1) determining and providing to another object the required coordinates in the terrain along with a respective resolution level; (2) receiving the data blocks corresponding to the specified coordinates; and (3) using the received data blocks to display a three-dimensional image. [*See* Ex. W; *see also* Ex. L at TERRABAS.HTM (describing data blocks with a Global Area Identifier corresponding to "coordinates in the terrain along with a respective resolution level," which allowed for "easy searching") and RENDERER.HTM (the renderer-database-manager "requests the data for a special location with an appropriate resolution"); Ex. I ('897 patent at col. 7:34-44 ("node 3 determines the field of view" and "calls up the data via the interchange network" after "selection of the earth as an object and input of a location and a direction of view"; required data is transmitted to node 3, which "sends this transmission for viewing")); Ex. J at GOOG 21376 (same).]

125. The preamble of claim 1 further recites "data blocks belonging to a hierarchical structure which includes blocks at a plurality of different resolution levels." The T_Vision system has "data blocks belonging to a hierarchical structure which includes blocks at a plurality of different resolution levels," as that phrase has been construed by the Court. [*See* Ex. W.] The T_Vision system has a "multi-layered database" that "is organized as a quadtree, containing higher levels of detail as you descend down the tree." [Ex. L at RENDERER.HTM, p. 1 of 2.]

Moreover, "the higher the number the finer the resolution." [*Id.* at TERRABAS.HTM, p. 1 of 3.] The T_Vision Project materials further describe this process of sub-dividing the planet:

> **Subdividing the Planet**
>
> Main part of the key is a string we call "Global Area Identifier" (GAI). Every patch has bounds generated by binary subdivision of the whole coordinate system:
>
> - longitude: (-180, 180) west <-> east
> - latitude : (90, -90) north <-> south
>
> There are three ways how this rectangle can be divided into 2 or 4 sectors, resulting in 8 possible sectors:

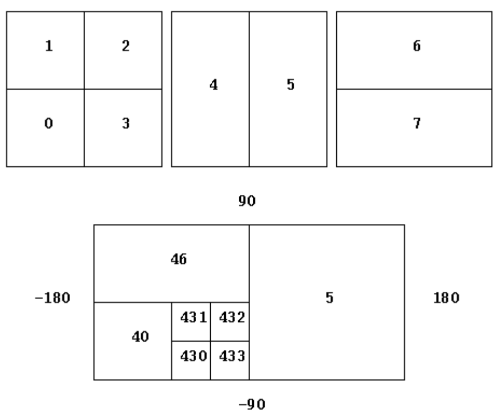

> The GAI can be seen as a kind of telephone number for reaching a particular sector of the planet. The number of digits corresponds to the level of detail; the higher the number, the finer the resolution. The reason for allowing horizontal and vertical subdivision and not only the four quadrants is that if you map the rectangle onto a sphere you can choose a subdivision strateg[y] that gives you more "quadratic" patches, especially around the poles.

[Ex. L at TERRABAS.HTM, p. 1-2 of 3; *see also* Ex. I ('897 patent at Fig. 5 (same Figure)); Ex. J at GOOG 21385.] Furthermore, the Mayer patent explains that the "sub-division of the image into sections with different spatial resolutions is preferably effected according to the method of a binary or quadrant tree." [Ex. I at ('897 patent at Abstract); *see also* Ex. J at GOOG 21372.] As elaborated in the Mayer patent, "The method according to the invention leads to a situation in which the data for the field of view to be shown are called up from the spatially distributed data sources only in the accuracy necessary for representation of the field of view with the desired image resolution, i.e. for example with high spatial resolution for close areas of the field of view or in low spatial resolution in a view to the horizon of a spherical object." [Ex. I ('897 patent at col. 2:62-3:2).] Correspondingly, as elaborated in the Mayer German patent application, "[t]he inventive method leads to a situation in which the data for the field of view to be represented are retrieved from the spatially distributed data sources only in the accuracy necessary for representing the field of view with the desired image resolution, i.e. for example with high spatial resolution for close areas of the field of view or in low spatial resolution in a view to the horizon of a spherical object." [Ex. J at GOOG 21374.] This is illustrated in Figure 6 of both

the Mayer patent and the Mayer German patent application, one of many figures in both prior art references showing data blocks belonging to a hierarchical structure. [Ex. I (‘897 patent at Fig. 6); Ex. J at GOOG 21386.] As explained in the Mayer patent: “FIG. 6 shows a sub-division according to an octant tree for a representation based on a three-dimensional geometrical model. Here a section 14 o[f] a space is sub-divided into eight spatial sub-sections 15. By means of the method according to the invention, consequently here also the data of just the spatial areas are called up in a higher accuracy, at which it is required in order to obtain the desired image resolution.” [Ex. I (‘897 patent at col. 8:55-62).] Correspondingly, as explained in the Mayer German patent application: “FIG. 6 shows a sub-division according to an octant tree for a representation based on a three-dimensional geometrical model. Here a section 14 of a space is sub-divided into eight spatial sub-sections 15. Consequently here also using the inventive method, the data of just those spatial areas are retrieved in a higher precision, in which this precision is required for obtaining the desired image resolution.” [Ex. J at GOOG 21377.]

c. *Claim 1:* receiving from the renderer one or more coordinates in the terrain along with indication of a respective resolution level

126. T_Vision uses “coordinates in the terrain” as construed by the Court, as stated in the Mayer patent: “a two-dimensional polygon grid model is used to display the data, which serves as a two-dimensional co-ordinate system for positioning the data. There were used as data to be displayed, for example satellite images, i.e. information relating to the colouring of the earth surface or geopolitical data or actual or stored meteorological data.” [Ex. I (‘897 patent at col. 7:10-16).] Analogously, in the Mayer German patent application: “a two-dimensional polygon grid model, which serves as a two-dimensional coordinate system for positioning the data, is used for representing the data. For example, satellite images, i.e. information about the coloring of the earth’s surface or geopolitical data or actual or stored meteorological data were used as data to be represented.” [Ex. J at GOOG 21376.] The T_Vision Project materials expressly reference data blocks having coordinates based on latitude and longitude (a Global Area Identifier (“GAI”)) and tout that this allows for “easy searching.” [Ex. L at TERRABAS.HTM, p. 1 of 3.] In addition, they note that “[t]he number of digits corresponds to the level of detail; the higher the number, the finer the resolution,” which is an indication of respective resolution level. [*Id; see also* Ex. I (‘897 patent at col. 8:35-38 & Fig. 4); Ex. J at GOOG 0021377.]

127. The T_Vision client computer is referred to as a node computer. As stated in the Mayer patent: “Systems of the company Silicon Graphics (SGI Onyx) were used as a node computer. This computer is capable of displaying more than 5,[0]00,000 texturised triangles per second and consequently is suitable for rapid picture build-up.” [Ex. I (‘897 patent at col. 6:22-25).] Correspondingly, as stated in the Mayer German patent application: “Systems of the company Silicon Graphics (SGI Onyx) were used as a node computer. This computer is capable of representing more than 5,[0]00,000 texturized triangles per second and is consequently suitable for rapid image formation.” [Ex. J at GOOG 9078.]

128. The T_Vision system practices the step of “receiving from the renderer one or more coordinates in the terrain along with indication of a respective resolution level.” [*See* Ex. W.] The T_Vision Project materials expressly disclose that “[t]he renderer computes the GAIs

according to the field of view and makes a simple query." [Ex. L at TERRABAS.HTM at p. 2 of 3.] Thus, the "coordinates in the terrain along with indication of respective resolution level" (i.e., GAIs) are sent by the renderer, and (by implication), received by another object. The T_Vision Project materials further provide that the renderer-database-manager "requests the data for a special location with an appropriate resolution." [Ex. L at RENDERER.HTM at p. 1 of 2.] In addition, the Mayer patent provides that "node 3 determines the field of view of the observer and calls up the data via the interchange network 7 and the nodes 1 and 2. These nodes in turn call up, via the collecting network 6, from the spatially distributed data sources 4 or for example from the camera 9, the required data and transmit them over the interchange network 7 to the node 3 for central storage." [Ex. I ('897 patent at col. 7:35-42).] Correspondingly, in the Mayer German patent application: "node 3 determines the field of view of the observer and retrieves the data using the exchange network 7 and the nodes 1 and 2. These nodes in turn retrieve the required data using the collective network 6 from the spatially distributed data sources 4 or, for example, from the camera 9, and transmit them using the exchange network 7 to the node 3 for central storage." [Ex. J at GOOG 21376.]

129. Furthermore, the "sections" referenced by the Mayer patent and the Mayer German patent application were requested with an indication of a respective resolution level. As described in the Mayer patent: "Fig. 4 shows the formation of an address of a section using the model of a quadrant tree for sub-division of the field of view 11. In the first sub-division of the field of view 11 into four sections 12, these are identified clockwise with the numerals 0 to 3. If a section is further sub-divided, the individual sub-sections 13 are numbered in the same way and the numbers thus obtained are prefixed to the numbers of the master section. With a permanently identical resolution of for example 128x128 points per section, the number of points of the section number is at the same time an indication of the level of spatial precision of the data." [Ex. I ('897 patent at col. 8:28-38).] Furthermore, the Mayer patent states, "Here a section 14 o[f] a space is sub-divided into eight spatial sub-sections 15. By means of the method according to the invention, consequently here also the data of just the spatial areas are called up in a higher accuracy, at which it is required in order to obtain the desired image resolution." [*Id.* ('897 patent at col. 8:57-62).] Correspondingly, as described in the Mayer German patent application: "FIG. 4 illustrates the formation of an address of a section using the model of a quadrant tree for sub-division of the field of view 11. In the first sub-division of the field of view 11 into four sections 12, these sections are identified clockwise with the numerals 0 to 3. If a section is further sub-divided, the individual sub-sections 13 are numbered in the same way and the number of the master section is prefixed to the number thus obtained." [Ex. J. at GOOG 21377.] Furthermore, the Mayer German patent application states, "Here a section 14 of a space is sub-divided into eight spatial sub-sections 15. Consequently here also using the inventive method, the data of just those spatial areas are retrieved in a higher precision, in which this precision is required for obtaining the desired image resolution." [*Id.*]

c. ***Claim 1:*** **providing the renderer with a first data block which includes data corresponding to the one or more coordinates, from a local memory**

130. The T_Vision system discloses "providing the renderer with a first data block which includes data corresponding to the one or more coordinates, from a local memory." [*See* Ex. W.] The T_Vision Project materials published at SIGGRAPH '95 provide, for example, that

the renderer "requests the data for a special location with an appropriate resolution…. If you approach too fast you will get a coarse image, but the frame rate is not affected." [Ex. L at RENDERER.HTM, p. 1 of 2.] I would understand from this disclosure that the "coarse" image is rendered from a first data block because the desired data block has not yet been downloaded. [*See also* Ex. L at TVISION.HTL, p. 2 of 3.] The Mayer patent states: "After each transmission and central storage of data, an image representation results, even if the data are insufficient to make possible the desired image resolution. As a result, even if the method is interrupted due to an alteration in the field of view and newly begun for a new field of view, the data for an image, even at low resolution, are always available." [Ex. I ('897 patent at col. 3:27-33).] The Mayer patent further discloses that, "The node 3 determines the representation of the data centrally stored therein and sends this…for viewing…to the display device 5," where "centrally stored" means locally stored. [Ex. I ('897 patent at col. 7:42-44).] Correspondingly, the Mayer German patent application discloses that, "The node 3 determines the representation of the data centrally stored therein and sends this transmission for representation to the display device 5." [Ex. J at GOOG 21376.] Correspondingly, the Mayer German patent application further states: "After each transmission and central storage of data, an image representation is carried out, even if the data are insufficient to enable the desired image resolution. As a result, even if the method is interrupted due to an alteration of the field of view and restarted for a new field of view, the data for an image, even at low resolution, are always available." [Ex. J at GOOG 21374.]

d. *Claim 1:* downloading from a remote server one or more additional data blocks at a resolution level higher than the resolution level of the first data block which include data corresponding to the one or more coordinates if the provided block from the local memory is not at the indicated resolution level.

131. The T_Vision system discloses "downloading from a remote server one or more additional data blocks at a resolution level higher than the resolution level of the first data block which include data corresponding to the one or more coordinates if the provided block from the local memory is not at the indicated resolution level." [*See* Ex. W.] The Mayer patent provides that, "[i]f the resolution of the representation is below the desired image resolution, the field of view is divided into sections and an investigation is undertaken for each individual section to see whether the data within the section are sufficient for a representation with the desired image resolution. If this is not the case for one of the sections, further data with a finer resolution are called up, transmitted and centrally stored from at least one of the spatially distributed data sources, and the section is shown with the new data. In turn an investigation is carried out into sufficient image resolution and possibly a further sub-division of the tested section is carried out into further partial sections as described above." [Ex. I ('897 patent at col. 2:17-29); *see also id.* at col. 7:45-59.] Correspondingly, the Mayer German patent application provides that: "Should the resolution of the representation be below the desired image resolution, the field of view is divided into sections and each individual section is examined to see whether the data within the section are sufficient for a representation with the desired image resolution. If this is not the case for any of the sections, further data with a finer resolution are retrieved for this section, transmitted and centrally stored from at least one of the spatially distributed data sources, and the section is represented with the new data. In turn an examination for sufficient image resolution and possibly a further sub-division of the examined section into further sub-sections are carried

out as described above." [Ex. J at GOOG 21373.] The subdividing discussed is inherent in using a hierarchical data structure from which to retrieve increasingly higher resolution image segments: "There was used, as a model for sub-dividing the field of view into sections and of these sections into further sections, a quadrant tree in which a progressive sub-division of an area into respectively four sections is carried out." [Ex. I ('897 patent at col. 7:30-33).] Correspondingly, in the Mayer German patent application, "A quadrant tree, in which a progressive sub-division of an area into four sections each is carried out, was used as a model for sub-dividing the field of view into sections and for sub-dividing these sections into further sections." [Ex. J at GOOG 21376.] Thus, the Mayer patent discloses a link between the condition of the provided block from local memory not being at the indicated resolution level and downloading additional data blocks. Moreover, "coarse" data blocks are downloaded first, and then progressively higher resolution data blocks are downloaded in the T_Vision system.

e. ***Claim 2:*** **A method according to claim 1, wherein downloading the one or more additional data blocks comprises downloading the blocks from a succession of resolution levels, from the level immediately higher than the resolution level of the first block up to the maximal existent resolution level on the server not above the indicated resolution level.**

132. The T_Vision system requests sections of the field of view in a coarse-to-fine order (i.e., from a level immediately higher than the resolution level of the first block up to the maximal existent resolution level on the server not above the indicated resolution level). In particular, the field of view is divided into four sections and, for each section having a resolution below the desired resolution, higher resolution space-related data is requested. This process is repeated until every section has the desired image resolution or no higher resolution data is available. In addition, the Mayer patent and the Mayer German patent application disclose transmitting data blocks either synchronously or asynchronously. For example, claim 1 of the Mayer patent encompasses transmitting data blocks either synchronously or asynchronously. This is illustrated by claim 41, which is a dependent claim to claim 1, and directed specifically to asynchronous transmission. During synchronous operation, when a data block is requested, its data is received before another block is requested, thus ensuring that data blocks will be received in the same order they are requested. Moreover, even during asynchronous operation, on average, data blocks will be received in the same order they are requested; however, there is no guarantee that they will be received in that order. Accordingly, based on my understanding of the Court's construction, the Mayer patent and the Mayer German patent application disclose this limitation of claim 2 at least to the extent they encompasses synchronous operation. [*See* Ex. W.] T_Vision also at least renders obvious this limitation. [*See id.*] In particular, T_Vision already suggests requesting sections of the field of view in order of increasing resolution level. Moreover, to the extent that Skyline seeks to construe claim 2 more broadly in an attempt to capture Google Earth, it is my opinion that under any broader construction, T_Vision would disclose this limitation of claim 2 whether it transmits data blocks synchronously or asynchronously.

f. *Claim 3*

133. Claim 3 is an independent claim. However, the limitations of claim 3 are substantially identical to limitations of claim 1. I address the differences below.

134. Claim 3 includes further language in the "receiving" step referring to the "plurality of coordinates being included in a plurality of respective distinct blocks." As discussed above with respect to claim 1, the renderer in T_Vision uses a two-dimensional coordinate system for positioning the data and a desired image resolution to call for additional data blocks. A Global Area Identifier, for example, designates the latitude, longitude and resolution level of each block. This falls within the Court's construction of claim 3 in my opinion. [*See* Ex. W.]

135. There are some minor variations in the language of the "providing" step of claim 3. However, these differences do not change my analysis: the T_Vision system practices this step for the same reasons discussed above with respect to claim 1.

136. Claim 3 also includes further language in the "downloading" step that "blocks of lower resolution levels are downloaded before blocks of higher resolution levels." The T_Vision system requests sections of the field of view in a coarse-to-fine order (i.e., blocks of lower resolution levels are downloaded before blocks of higher resolution levels). In particular, the field of view is divided into four sections and, for each section having a resolution below the desired resolution, higher resolution space-related data is requested. This process is repeated until every section has the desired image resolution or no higher resolution data is available. In addition, the Mayer patent and the Mayer German patent application disclose transmitting data blocks either synchronously or asynchronously. For example, claim 1 of the Mayer patent encompasses transmitting data blocks either synchronously or asynchronously. This is illustrated by claim 41, which is a dependent claim to claim 1, directed specifically to asynchronous transmission. During synchronous operation, when a data block is requested, its data is received before another block is requested, thus ensuring that data blocks will be received in the same order they are requested. Moreover, even during asynchronous operation, on average, data blocks will be received in the same order they are requested; however, there is no guarantee that they will be received in that order. Accordingly, based on my understanding of the Court's construction, the Mayer patent and the Mayer German patent application disclose this limitation of claim 3 at least to the extent they encompasses synchronous operation. [*See* Ex. W.] T_Vision also at least renders obvious this limitation. [*See id.*] In particular, T_Vision already suggests requesting sections of the field of view in order of increasing resolution level. Moreover, to the extent that Skyline seeks to construe claim 3 more broadly in an attempt to capture Google Earth, it is my opinion that under any broader construction, T_Vision would disclose this limitation of claim 3 whether it transmits data blocks synchronously or asynchronously.

g. *Claim 5*

137. Claim 5 is an independent claim. However, the limitations of claim 5 are substantially identical to limitations of claim 1 and 3. I address the differences below.

138. Claim 5 includes further language in the "downloading" step that "the blocks are downloaded according to the order in which the coordinates were provided." The T_Vision system requests sections of the field of view in a coarse-to-fine order (i.e., the blocks are downloaded according to the order in which the coordinates were provided). In particular, the field of view is divided into four sections and, for each section having a resolution below the desired resolution, higher resolution space-related data is requested. This process is repeated until every section has the desired image resolution or no higher resolution data is available. In addition, the Mayer patent and the Mayer German patent application disclose transmitting data blocks either synchronously or asynchronously. For example, claim 1 of the Mayer patent encompasses transmitting data blocks either synchronously or asynchronously. This is illustrated by claim 41, which is a dependent claim to claim 1, directed specifically to asynchronous transmission. During synchronous operation, when a data block is requested, its data is received before another block is requested, thus ensuring that data blocks will be received in the same order they are requested. Moreover, even during asynchronous operation, on average, data blocks will be received in the same order they are requested; however, there is no guarantee that they will be received in that order. Accordingly, based on my understanding of the Court's construction, the Mayer patent and the Mayer German patent application disclose this limitation of claim 5 at least to the extent they encompasses synchronous operation. [*See* Ex. W.] T_Vision also at least renders obvious this limitation. [*See id.*] In particular, T_Vision already suggests requesting sections of the field of view in order of increasing resolution level. Moreover, to the extent that Skyline seeks to construe claim 5 more broadly in an attempt to capture Google Earth, it is my opinion that under any broader construction, T_Vision would disclose this limitation of claim 5 whether it transmits data blocks synchronously or asynchronously.

h. *Claim 7*

139. Claim 7 is an independent claim. However, the limitations of claim 7 are substantially identical to limitations of claim 1. I address the differences below.

140. Claim 7 includes further language in the "downloading" step requiring "downloading from a remote server excess blocks not currently needed by the renderer to fill up the local memory when not downloading blocks required by the renderer." Both the Mayer patent and Mayer German patent application disclose a separate process for downloading excess data blocks from the regions surrounding the field of view. [*See* Ex. W.] This process determines the probability that an area will pass within the field of view, and, based on this determination, downloads data blocks from the areas with the highest probability. These references do not explicitly state that this process only runs during periods of time when the local computer is not downloading data blocks needed by the renderer. However, at minimum, it would have been obvious to a person of ordinary skill in the art to first obtain data blocks needed to render the current view before downloading excess data blocks, and, in my opinion, this disclosure is actually inherent in the Mayer patent and Mayer German patent application.

i. *Claim 8:* A method according to claim 7, wherein downloading the data blocks comprised downloading the blocks via the Internet.

141. Claim 8 includes all the limitations of claim 7, and a further limitation relating to "downloading the blocks via the Internet." T_Vision was recommended for use with an asynchronous transmission protocol with a transmission rate greater than 35 MBits/second. The T_Vision system was implemented and demonstrated using a high-speed ATM network. This was not a solely private network, but was part of the broader Internet, accessible, for example, over the Internet from SIGGRAPH '95. [*See* Clinger Decl. & Exs.] In my opinion, use of the T_Vision network over the broader Internet would fall within the Court's construction. Moreover, it would have been obvious to download data blocks in T_Vision via any network that met the recited transmission rate, including the Internet.

142. Finally, there was a trend in the art to adapt 3D terrain visualization applications for use on the Internet, including the World Wide Web. This is illustrated, for example, by Rabinovich: "The advent of the World-Wide-Web suggests the running of this type of [terrain visualization] application over the Internet, in a client/server scenario." [Ex. Q at GOOG 9828.] Likewise, Potmesil describes a "WWW-based system … which allows users to view, search and post-geographically-indexed information of the Earth," including "a 3D flight-simulator browser capable of continuous flight around the Earth." [Ex. R at GOOG 9762.] He states that "Internet-based computers and communications can be very effective in enhancing our ability to visualize and to search 3D environments in the great outdoors where we move, work, play and learn." [*Id.*] One of ordinary skill in the art would have been motivated to combine the World Wide Web-based strategies of Rabinovich and Potmesil with ATM-based terrain visualization applications such as TerraVision and T_Vision. Potmesil specifically references ATM-based terrain visualization applications, including T_Vision. [*Id.* at GOOG 9764.] This combination is also suggested by the TerraVision system, as TerraVision was also demonstrated over the Internet despite being designed for a high-speed ATM network. Accordingly, in my opinion, claim 8 would have at least been rendered obvious by the combination of T_Vision with one of Rabinovich, Potmesil or TerraVision.

j. *Claim 9:* A method according to claim 7, wherein the renderer renders a view from a current viewpoint, and wherein downloading the excess blocks comprises filling the local memory with substantially all of the blocks surrounding a point in the terrain seen from the current viewpoint within a predetermined distance range.

143. Claim 9 includes all the limitations of claim 7, and a further limitation relating to downloading "substantially all the blocks surrounding a point in the terrain seen from the current viewpoint within a predetermined range." As discussed above with respect to claim 7, the Mayer patent and Mayer German patent application disclose a separate process for downloading excess data blocks from the areas with the highest probability of passing within the field of view. In many typical cases, use of this technique will result in substantially all of the blocks covering terrain within a uniform predetermined distance from a point in the terrain that is seen from the current viewpoint being downloaded.

144. Furthermore, downloading blocks surrounding a point in the terrain seen from the current viewpoint within a predetermined distance range was known in the art. This is shown, for example, by Potmesil, who discloses the downloading of excess data blocks using several different caching strategies, one of which is downloading blocks surrounding a point in the terrain seen from the current viewpoint within a predetermined distance range was known in the art. [*See* Ex. R at GOOG 9767.] He states: "The caching algorithm uses the user's current position, velocity, and acceleration to estimate where the user is moving and allocates new tiles there." [*Id.*] One of the "several implemented caching strategies" is "obtain as many tiles as close to the view as possible; this may be used in low speeds (hovering) when the direction of flight is uncertain." [*Id.*] Potmesil also discloses filling the local memory with excess tiles: "The more disc and memory space the host machine has available, the more tiles can be brought into the cache and remain there." [*Id.*] One of ordinary skill in the art would have been motivated to combine the caching strategies of Potmesil with T_Vision as Potmesil specifically references T_Vision (another terrain visualization application). [*Id.* at GOOG 9764.] Accordingly, in my opinion, claim 9 would have at least been rendered obvious by the combination of one of TerraVision or T_Vision with Potmesil.

k. *Claim 11:* A method according to claim 9, wherein filling the local memory comprises filling the memory with substantially all the blocks within the range from a lower resolution level before downloading blocks of higher resolution levels.

145. Claim 11 includes all the limitations of claims 7 and 9, and a further limitation relating to filling the local memory with "substantially all the blocks within the range from a lower resolution level before downloading blocks of higher resolution level." The T_Vision system requests blocks in a coarse-to-fine order (i.e., blocks of lower resolution levels are downloaded before blocks of higher resolution levels). In particular, the field of view is divided into four sections and, for each section having a resolution below the desired resolution, higher resolution space-related data is requested. This process is repeated until every section has the desired image resolution or no higher resolution data is available. In addition, the Mayer patent and the Mayer German patent application disclose transmitting data blocks either synchronously or asynchronously. For example, claim 1 of the Mayer patent encompasses transmitting data blocks either synchronously or asynchronously. This is illustrated by claim 41, which is a dependent claim to claim 1, directed specifically to asynchronous transmission. During synchronous operation, when a data block is requested, its data is received before another block is requested, thus ensuring that data blocks will be received in the same order they are requested. Moreover, even during asynchronous operation, on average, data blocks will be received in the same order they are requested; however, there is no guarantee that they will be received in that order. Accordingly, based on my understanding of the Court's construction, the Mayer patent and the Mayer German patent application disclose this limitation of claim 11 at least to the extent they encompasses synchronous operation. [*See* Ex. W.] T_Vision also at least renders obvious this limitation. [*See id.*] In particular, T_Vision already suggests requesting sections of the field of view in order of increasing resolution level. Moreover, to the extent that Skyline seeks to construe claim 11 more broadly in an attempt to capture Google Earth, it is my opinion that under any broader construction, T_Vision would disclose this limitation of claim 11 whether it transmits data blocks synchronously or asynchronously.

k. *Claim 12:* An apparatus [of claim 1]

146. The limitations of claim 12 are substantially identical to those in claim 1, except that claim 12 is directed to an "apparatus" rather than a "method."

147. The T_Vision system was software designed to be run on a computer (i.e., an apparatus). "Systems of the company Silicon Graphics (SGI Onyx) were used as a node computer." [Ex. I ('897 patent at col. 6:22-23); Ex. J at GOOG 21376.] Thus, for the reasons detailed above with respect to claim 1, the T_Vision system also discloses every limitation of claim 12. Moreover, as discussed in more detail above with respect to TerraVision, there is nothing in claim 12 that would somehow read out an SGI Onyx computer. Claim 12 refers to an "apparatus" with a "local memory," a "communications link" and a "processor"—elements found in the SGI workstation used in the T_Vision system. T_Vision refers to the "local memory" as a "central data memory" and the communications link as a "data transmission network."

148. In addition, the mid 1990s SGI workstations described in the T_Vision project, the Mayer patent and the Mayer German patent application (SGI Onyx), were comparable to or less powerful than many off-the-shelf PCs running Microsoft Windows available by February 1999, when the '189 patent was filed.

l. *Claim 13:* An apparatus [of claim 2]

149. The limitations of claim 13 are substantially identical to those of claims 1 and 2, except that claim 13 is directed to an "apparatus" rather than a "method." For the reasons stated with respect to claims 1, 2, and 12, in my opinion, T_Vision anticipates and/or renders obvious this claim.

m. *Claim 14:* An apparatus [of claim 3]

150. The limitations of claim 14 are substantially identical to those of claim 3, except that claim 14 is directed to an "apparatus" rather than a "method." For the reasons stated with respect to claims 3 and 12, in my opinion, T_Vision anticipates and/or renders obvious this claim.

n. *Claim 16:* An apparatus [of claim 5]

151. The limitations of claim 16 are substantially identical to those of claim 5, except that claim 16 is directed to an "apparatus" rather than a "method." For the reasons stated with respect to claims 5 and 12, in my opinion, T_Vision at least renders obvious this claim.

o. *Claim 18:* An apparatus [of claim 7]

152. The limitations of claim 18 are substantially identical to those of claim 7, except that claim 18 is directed to an "apparatus" rather than a "method." For the reasons stated with respect to claims 7 and 12, in my opinion, T_Vision anticipates and/or renders obvious this claim.

p. *Claim 19:* **An apparatus [of claim 9]**

153. The limitations of claim 19 are substantially identical to those of claims 7 and 9, except that claim 19 is directed to an "apparatus" rather than a "method." For the reasons stated with respect to claims 7, 9, and 12, in my opinion, T_Vision anticipates and/or renders obvious this claim.

q. *Claim 21:* **An apparatus [of claim 11]**

154. The limitations of claim 21 are substantially identical to those of claims 7, 9 and 11, except that claim 21 is directed to an "apparatus" rather than a "method." For the reasons stated with respect to claims 7, 9, 11, and 12, in my opinion, T_Vision anticipates and/or renders obvious this claim.

r. *Claim 22:* **An apparatus [of claim 8]**

155. The limitations of claim 22 are substantially identical to those of claims 7 and 8, except that claim 22 is directed to an "apparatus" rather than a "method." Furthermore, the Court has construed this claim as only requiring a "connection to the Internet" as opposed to requiring that "some portion of the download [] happen over the broader Internet." [*See* Ex. D-2 at 22.] As discussed with respect to claim 8, T_Vision was at least connected to the broader Internet and/or it would have been obvious to connect T_Vision to the broader Internet. For the reasons stated with respect to claims 7, 8, and 12, in my opinion, T_Vision anticipates and/or renders obvious this claim.

s. *Claim 23:* **The method of claim 7, wherein the coordinates relate to the coordinates of a predetermined course of a flight vehicle.**

156. Claim 23 includes all the limitations of claim 7, and a further limitation providing that "the coordinates relate to the coordinates of a predetermined course of a flight vehicle." The T_Vision system performs this further limitation of claim 7. [*See* Ex. W.] In particular, the Mayer patent and Mayer German patent application disclose that the T_Vision system is an improved flight simulator. Indeed, Potmesil refers to T_Vision as a "noteworthy 3D flight simulator." [*See* Ex. R at GOOG 9764.] The T_Vision system also allowed the user to roam about the terrain at will.

157. Furthermore, there was a trend in the art to use terrain visualization systems for flight simulation. The Migdal patent is evidence of this trend: "As is well-known in graphics design, the display view can simulate flight by constantly moving the eyepoint along a terrain or landscape being viewed. Such flight can be performed automatically as part of a program application, or manually in response to user input such as a mouse or joystick movement." [Ex. N ('783 patent at col. 10:16-21).] A person of ordinary skill in the art would have been motivated to apply this well-known trend to other terrain visualization systems such as TerraVision and T_Vision. Accordingly, in my opinion, claim 23 would have at least been rendered obvious by the combination of one of TerraVision or T_Vision with the Migdal patent.

t. *Claim 24:* An apparatus [of claim 23]

158. The limitations of claim 24 are substantially identical to those of claims 7 and 23, except that claim 24 is directed to an "apparatus" rather than a "method." For the reasons stated with respect to claims 7, 12, and 23, in my opinion, T_Vision anticipates and/or renders obvious this claim.

3. *Response to Manocha Report re: T_Vision*

159. With respect to T_Vision, Dr. Manocha takes the same approach he did with TerraVision, attempting to distinguish this reference based on its "objective and approach" rather than the claim language as construed by the Court. T_Vision was a software application designed for use on a high-speed network. A prototype of this system was demonstrated at SIGGRAPH '95. Moreover, the Mayer patent and the Mayer German patent application also make further disclosures of a terrain visualization system as detailed in Exhibit W.

160. Dr. Manocha suggests that none of the T_Vision references discloses the use of three-dimensional terrain data. However, all of these references refer to both image and elevation data, not just high resolution satellite images. Indeed, the T_Vision Project materials use the term "DEM data," which Dr. Manocha acknowledged in connection with TerraVision was terrain data. These disclosures are more than sufficient to show that T_Vision was capable of rendering three-dimensional terrain. Moreover, these references do not distinguish between downloading image data and downloading elevation data but rather disclose downloading both in exactly the same way—in coarse-to-fine order as needed. That T_Vision rendered three-dimensional terrain data is also graphically demonstrated and described in the T_Vision Video:

> By switching the surface off, we can observe this process more easily. As the distance between us and the Earth increases, the high resolution data is removed from the memory and is replaced with new data for the wider field of view. Out of these different levels of altitude data, we compute the tectonic surface of the Earth and then project the corresponding satellite images onto it. An asynchronous and anticipatory loading strategy always guarantees a steady frames per second.

[Ex. M.]. In addition to the regular textured view, this portion of the T_Vision Video shows both wireframe and shaded, untextured views of the terrain that clearly indicate that it was rendered from a 3D model whose elevation varies across the terrain:

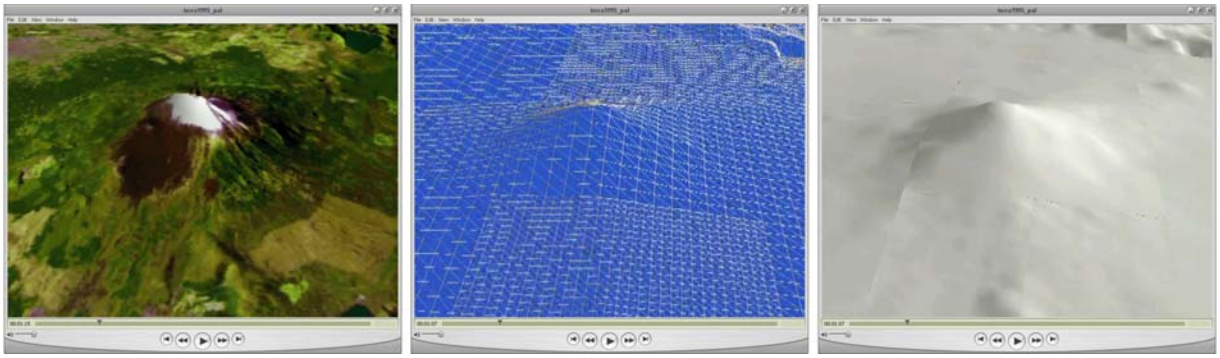

161. Dr. Manocha states that the Mayer patent and the other references do not disclose the claimed method of using a first data block from local memory. That is not correct. Initially, all of the terrain data in the T_Vision system is stored remotely. However, a "first data set" is stored in local memory and displayed, and the T_Vision system will then download additional data blocks as needed.

162. The fact that the references describing the T_Vision system did not use the term "data block" is also of no import. Data in T_Vision is stored in multi-resolution hierarchies, organized as "quadtree[s], containing higher levels of detail as you descend the tree." [Ex. L at RENDERER.HTM; *see also* Ex. I ('897 patent at col. 8:28-42 & Fig. 4); Ex. J at GOOG 21377, -84.] Sections of data in the hierarchy—i.e., one or more data blocks—are downloaded as needed in T_Vision. [*Id.*]

163. Dr. Manocha also again misunderstands my opinions regarding the "receiving from" and "providing to" elements of claim 1. The renderer in T_Vision determines the coordinates in the terrain and provides these coordinates to another object for the purpose of downloading additional data blocks from the remote server if the data blocks in local memory are not at the requested resolution level.

164. Dr. Manocha further suggests that the T_Vision system demonstrated at SIGGRAPH '95 was not actually connected to an ATM network. However, an ATM network was set up for the SIGGRAPH '95 conference and it is my understanding that T_Vision had a connection to that network and the broader Internet. [Clinger Decl. & Exs.] Furthermore, even assuming that the T_Vision system at SIGGRAPH '95 was not actually connected to a network, it is clear from the materials presented at SIGGRAPH '95 that T_Vision was designed for use on such a network. [*See* Exs. L & K.]

165. Finally, Dr. Manocha appears to ignore my statements regarding anticipation by the Mayer patent and the Mayer German patent application and claims that I am only relying on the T_Vision system for anticipation. That is not the case. I analyzed and identified relevant disclosures in each of the T_Vision references on an element-by-element basis. However, for convenience, I have attached as Exhibit W to this report a chart setting out these disclosures.

4. ***Conclusions***

166. Based on my comparison of the asserted claims to the T_Vision system and to the individual T_Vision references, it is my opinion that the T_Vision system and references anticipate or render obvious all asserted claims of the '189 patent. Specifically, in my opinion, claims 1-3, 5, 7-9, 11-14, 16, 18-19, and 21-24 are anticipated and/or rendered obvious by the T_Vision system, the Mayer patent and the Mayer German patent application.

167. Further, a person of ordinary skill in the art would have been motivated to combine any of the T_Vision references, since they all refer to exactly the same project and system. Furthermore, any alleged differences between asserted claims and the prior art are such that the claims would have been obvious at the time the invention was made to a person having ordinary skill in the art. In particular, the T_Vision achieved all of the objectives of the claimed inventions and any differences would have merely related to implementation details.

168. In my opinion, the combination of one or more of the T_Vision system, the Mayer patent, and the Mayer German patent application, with one or more of the TerraVision system, the MAGIC Final Report, the TerraVision Technical Note, and the MAGIC IEEE Article, also render the asserted claims at least obvious. The T_Vision system and the TerraVision system were demonstrated next to or across from each other at SIGGRAPH '95 in the Interactive Communities venue, and were both described on the SIGGRAPH '95 Multimedia CD-ROMs; accordingly, one of ordinary skill in the art would have been motivated to combine these references, which were both directed to systems for visualizing terrain and were implemented using similar software and hardware processes. Since a person of ordinary skill in the art would have been motivated to combine the systems together, he or she also would have been motivated to combines references about these systems.

E. Migdal and Cosman

169. In my opinion, the Migdal patent and Cosman in combination, render claims 1-3, 5, 12-14, and 16 of the '189 patent obvious.

1. *Overview of the Migdal patent and Cosman*

170. The Migdal patent (Ex. N) was filed in the United States on Nov. 6, 1995, and issued on Jun. 2, 1998. The Migdal International Application (Ex. O) was filed on November 6, 1995, and was published on May 15, 1997. Thus, it is my understanding that both the Migdal patent and the Migdal International Application constitute prior art with respect to the '189 patent. It is my understanding that the Migdal International Application corresponds to the United States patent application that resulted in the Migdal patent because the priority application of the Migdal International Application is this United States patent application.

171. Cosman (Ex. P) is a printed publication, published no later than July 1997. Thus, it is my understanding that Cosman constitutes prior art with respect to the '189 patent.

2. *Comparing the Migdal patent and Cosman to the Asserted Claims of the '189 Patent*

172. In my opinion, the Migdal patent renders obvious claims 1-3, 5, 12-14, and 16 of the '189 patent in combination with Cosman.

173. In the paragraphs below, I have not repeated the Court's construction of relevant claim terms or opinions relating to the interpretation of the '189 patent, which are set forth in detail above with respect to TerraVision. However, I incorporate these discussions herein by reference as they are equally applicable.

a. *Claim 1, Preamble:* **A method of providing data blocks describing three-dimensional terrain data to a renderer, the data blocks belonging to a hierarchical structure which includes blocks at a plurality of different resolution levels**

174. The Migdal patent discloses "a method of providing data blocks describing three-dimensional terrain to a renderer." In the Abstract of the Migdal patent, the inventors state that

they are disclosing "an apparatus and method for … providing texel [texture element] data relevant for displaying a textured image." [Ex. N ('783 patent at Abstract).] Furthermore, the Migdal patent provides that "three-dimensional texture data can be used." [*Id.* ('783 patent at col. 3:18-19).] The Migdal patent also describes a processor (202) and graphics subsystem (220) which "maps texture data … to pixel data in the screen space." [*Id.* ('783 patent at col. 6:65-7:21).] That is, the texel data, or "data blocks," are provided to a "renderer." The "data blocks" can be "data blocks describing three-dimensional terrain data" because the texture data can represent terrain: "A large amount of texture source data, such as photographic terrain texture, is stored as a two-dimensional or three-dimensional texture MIP-map…. Only a relatively small clip-map representing selected portions of the complete texture MIP-map is loaded into faster, more expensive memory." [*Id.* ('783 patent at Abstract).]

175. The Migdal patent also discloses a "renderer" (at least as interpreted by Skyline), that is, software and/or hardware that performs the functions of (1) determining and providing to another object the required coordinates in the terrain along with a respective resolution level; (2) receiving the data blocks corresponding to the specified coordinates; and (3) using the received data blocks to display a three-dimensional image. In the Migdal patent, processor 202 determines the new tiles that are needed "when the eyepoint and/or field of view changes to ensure that clip-map 540 contains the texel data which is most likely to be rendered for display", processor 202 receives the new tiles from mass storage device 208 and "passes" them to texture memory 226, and graphics subsystem 220 renders an image based on the texture data in texture memory 226. [Ex. N ('793 patent at col. 10:42-52, col. 6:55-57, & col. 6:65-7:21, respectively).]

176. The Migdal patent further discloses "data blocks belonging to a hierarchical structure which includes blocks at a plurality of different resolution levels." [Ex. N ('793 patent at col. 1:47-49 ("The MIP-map consists of a texture pattern pre-filtered at progressively lower or coarser resolutions and stored in varying levels of detail (LOD) maps. See, e.g., the explanation of conventional texture MIP-mapping in Foley et al., *Computer Graphics Principles and Practice,* Second Edition, Addison-Wesley Publishing Company, Reading, Mass. (1990), pages 742–43 and 826–28 (incorporated by reference herein).")).] Particularly, the Migdal patent "pertains to an apparatus and method for providing texture by using selected portions of a texture MIP-map." [*Id.* ('793 patent at col. 3:6-9).]

b. *Claim 1:* receiving from the renderer one or more coordinates in the terrain along with indication of a respective resolution level

177. The Migdal patent specifically discloses such use of "coordinates in the terrain along with indication of a respective resolution level." First, the coordinates in the terrain are computed: "Texture coordinates for each pixel quad are ultimately output from scan conversion module 1030. Normalizer and divider 1040 outputs normalized texture coordinates for the pixel quad. Such scan conversion processing is well-known for both two-dimensional and three-dimensional texture mapping and need not be described in further detail." Ex. N ('783 patent at col. 14:28-33). Next, the respective resolution level is computed: "LOD generation block 1050 first calculates an appropriate level of detail for the pixel quad (or pixel) according to standard LOD generation methods based on the individual pixel size and texel dimension." Ex. N ('783 patent at col. 14:45-49.) The "appropriate level of detail" is the same thing as the "respective

resolution level." Thus, the system described in the Migdal patent receives from the "renderer" coordinates along with indication of a respective resolution level.

c. *Claim 1:* providing the renderer with a first data block which includes data corresponding to the one or more coordinates, from a local memory

178. The Migdal patent discloses "providing the renderer with a first data block which includes data corresponding to the one or more coordinates, from a local memory." The Migdal patent teaches: "According to the present invention, then, a check is made to determine whether texel data for the pixel quad is included within a tile corresponding to the appropriate level of detail. When a texel is included within a tile at the appropriate level of detail, a LOD value corresponding to this tile is output. Otherwise, a LOD value is output identifying a tile at a lower level of detail which includes a substitute texel." Ex. N ('783 patent at col. 14:45-58). Thus, the "renderer" is given access to the tile in a local memory that contains the desired texel corresponding to the coordinates, though possibly at a coarser level of the hierarchy (i.e., "a first data block which includes data corresponding to the one or more coordinates, from a local memory").

d. *Claim 1:* downloading from a remote server one or more additional data blocks at a resolution level higher than the resolution level of the first data block which include data corresponding to the one or more coordinates if the provided block from the local memory is not at the indicated resolution level.

179. The Migdal patent describes how the clip-map is updated by loading the clip-map with relevant portions (data blocks) of the entire MIP-map from a mass storage device. [*Id.* ('783 patent at cols. 10:55–11:23).] This updating can be conditional on changes in viewpoint: "For example, as the eyepoint shifts, a new row of texel data is added to a tile in the direction of the eyepoint movement. A row located away from a new eyepoint location is discarded. Coarser tiles need not be updated until the eyepoint has moved sufficiently far to require a new row of texel data." [*Id.* ('783 patent at col. 10:55–62).] Thus, the same change in viewpoint that necessitated the provided block not being at the indicated resolution, causes additional data blocks to be read into the clip-map. The Migdal patent further discloses that a mass storage device containing texture data may be remote: "[u]nder conventional texture mapping techniques, even if texture data were to be accessed from a remote, large texture MIP-map, the rendering of a textured image for display in real-time would be impractical, if not impossible." [Ex. N ('783 patent at col. 7:51-58).] Thus, the Migdal patent discloses the use of remote locations for storing MIP-map data. Furthermore, the Migdal patent discloses that its system may be connected to a network. [*See, e.g.*, *id.* ('783 patent at col. 7:18-21 (stating that the rendered displayed image may be transmitted over a network)).] Therefore, it would be obvious to someone of ordinary skill in the art to practice the teachings of the Migdal patent by downloading data blocks from a remote MIP-map accessed over a network when the MIP-map would be too large to store on the local machine. For example, the Migdal patent states, "This hierarchical texture mapping storage scheme allows huge texture MIP-maps to be stored rather inexpensively on the mass storage device 208." A person of ordinary skill would also be aware

that a mass storage device could be accessed through a network instead of being local to the machine.

180. Cosman discloses that "[t]errain texture is paged independently for each MIP LOD, and this provides a large reduction in total paging load. Preferential paging of the lower LODs ensures a benign paging overload fallback--terrain would then be textured with a lower resolution version of the right texture." [Ex. P (Cosman, p. 62 ¶ 2).] Thus, in Cosman, lower level of detail ("LOD") blocks are obtained first, so that if the system gets overloaded and cannot obtain the higher LOD blocks needed for an image in time, the terrain can still be textured with a lower resolution version of the correct texture.

181. The Migdal patent discloses loading a relatively small clip-map representing selected portions of a complete texture MIP-map into a fast memory for rapidly displaying an image. Attempts to access a texture element lying outside of a particular clip-map tile are accommodated by utilizing a substitute texture element obtained from the next coarser resolution clip-map tile that encompasses the desired texture element.

182. One of ordinary skill in the art would have been motivated to combine the technique of Cosman with the system of the Migdal patent in order to ensure that coarser levels of the clip-map are loaded before finer levels of the clip-map when multiple levels of the clip-map must be loaded. The suggestion or motivation for combining the prior art references identified above can be found in the references themselves: Cosman was a cited reference in the Migdal patent. In addition, the Cosman and Migdal references both relate to the field of computer graphics and are both directed to solving the problem of rapidly providing texture data for displaying an image. Further motivation to combine prior art references includes trends in the art, exercise of ordinary skill in the art, and knowledge generally available to those of ordinary skill in the art.

e. ***Claim 2:*** **A method according to claim 1, wherein downloading the one or more additional data blocks comprises downloading the blocks from a succession of resolution levels, from the level immediately higher than the resolution level of the first block up to the maximal existent resolution level on the server not above the indicated resolution level.**

183. Cosman performs this further limitation of claim 2. Cosman discloses that "[t]errain texture is paged independently for each MIP LOD, and this provides a large reduction in total paging load. Preferential paging of the lower LODs ensures a benign paging overload fallback--terrain would then be textured with a lower resolution version of the right texture." [Ex. P (Cosman, p. 62 ¶ 2).] In Cosman, lower LODs are obtained first, so that if the system gets overloaded and cannot obtain the higher LODs needed for an image in time, the terrain can still be textured with a lower resolution version of the correct texture. Thus, in my opinion, the Migdal patent in combination with Cosman renders obvious this claim.

f. *Claim 3*

184. Claim 3 is an independent claim. However, the limitations of claim 3 are substantially identical to limitations of claim 1. I address the differences below.

185. Claim 3 includes further language in the "receiving" step that "said plurality of coordinates being included in a plurality of respective distinct blocks." As discussed above with respect to claim 1, the renderer in the Migdal patent uses a scan conversion process to obtain a plurality of coordinates for the LOD data blocks. This falls within the Court's construction of claim 3 in my opinion.

186. There are some minor variations in the language of the "providing" step of claim 3. However, these differences do not change my analysis: the Migdal patent discloses this step for the same reasons discussed above with respect to claim 1.

187. Claim 3 also includes further language in the "downloading" step that "blocks of lower resolution levels are downloaded before blocks of higher resolution levels." Cosman discloses that "[t]errain texture is paged independently for each MIP LOD, and this provides a large reduction in total paging load. Preferential paging of the lower LODs ensures a benign paging overload fallback--terrain would then be textured with a lower resolution version of the right texture." [Ex. P (Cosman, p. 62 ¶ 2).] In Cosman, lower LODs are obtained first, so that if the system gets overloaded and cannot obtain the higher LODs needed for an image in time, the terrain can still be textured with a lower resolution version of the correct texture. Thus, in my opinion, the Migdal patent in combination with Cosman renders obvious this claim.

g. *Claim 5*

188. Claim 5 is an independent claim. However, the limitations of claim 5 are substantially identical to limitations of claim 1 and 3. I address the differences below.

189. Claim 5 includes further language in the "downloading" step that "the blocks are downloaded according to the order in which the coordinates were provided." Cosman discloses that "[t]errain texture is paged independently for each MIP LOD, and this provides a large reduction in total paging load. Preferential paging of the lower LODs ensures a benign paging overload fallback--terrain would then be textured with a lower resolution version of the right texture." [Ex. P (Cosman, p. 62 ¶ 2).] In Cosman, lower LODs are obtained first, so that if the system gets overloaded and cannot obtain the higher LODs needed for an image in time, the terrain can still be textured with a lower resolution version of the correct texture. A person of ordinary skill in the art would understand that if two pixels to be rendered in sequence both require texture data that is not in the clip-map, and each requires data from different MIP-map blocks, then after providing a first data block for one or both pixels, it would be obvious to download a block at the appropriate resolution for the first pixel before downloading a block at the appropriate resolution for the second pixel. Thus, in my opinion, the Migdal patent in combination with Cosman renders obvious this claim.

h. *Claim 12:* An apparatus [of claim 1]

190. The limitations of claim 12 are substantially identical to those in claim 1, except that claim 12 is directed to an "apparatus" rather than a "method."

191. The systems of both the Migdal patent and Cosman were designed to be run on a computer (i.e., an apparatus). Thus, for the reasons detailed above with respect to claim 1, the Migdal patent and Cosman in combination render claim 12 obvious. In my opinion, there is nothing in claim 12 that would somehow read out either of the systems disclosed in the Migdal patent and Cosman.

i. *Claim 13:* An apparatus [of claim 2]

192. The limitations of claim 13 are substantially identical to those of claims 1 and 2, except that claim 13 is directed to an "apparatus" rather than a "method." For the reasons stated with respect to claims 1, 2 and 12, in my opinion, the Migdal patent in combination with Cosman renders obvious this claim.

j. *Claim 14:* An apparatus [of claim 3]

193. The limitations of claim 14 are substantially identical to those of claim 3, except that claim 14 is directed to an "apparatus" rather than a "method." For the reasons stated with respect to claims 3 and 12, in my opinion, the Migdal patent in combination with Cosman renders obvious this claim.

k. *Claim 16:* An apparatus [of claim 5]

194. The limitations of claim 16 are substantially identical to those of claim 5, except that claim 16 is directed to an "apparatus" rather than a "method." For the reasons stated with respect to claims 5 and 12, in my opinion, the Migdal patent in combination with Cosman renders obvious this claim.

3. *Response to Manocha Report re: Migdal and Cosman*

195. With respect to Migdal and Cosman, Dr. Manocha cites testimony and documents from Michael Jones, and asserts that these establish that neither Migdal nor Cosman invalidates the '189 patent. However, in both the testimony and documents, Mr. Jones states that he did not perform an invalidity analysis or reach a definite opinion regarding whether Migdal or Cosman invalidates the claims of the '189 patent. Moreover, Mr. Jones does note that there is significant overlap between the prior art and the '189 patent. In pointing to Mr. Jones' testimony, Dr. Manocha also avoids analyzing the actual disclosures in the prior art that contradict his conclusions.

196. Dr. Manocha claims further that the PTO reviewed both Migdal and Cosman during its examination of the '189 patent. However, nowhere is Cosman referenced in the cited art for the '189 patent and there is no evidence that the Examiner ever considered this reference.

197. Moreover, while Dr. Manocha argues that the PTO's examination of the Migdal patent should be given deference, he ignores the Examiner's finding that Migdal discloses every limitation of the claims at issue except the download order—a limitation expressly found in the Cosman reference.

198. Dr. Manocha also misstates my opinions in certain respects, again claiming I have "admitted" that certain elements are missing in Migdal and Cosman.

199. Finally, Dr. Manocha's criticism that I have not put my report in the form of a chart is also unwarranted—I analyzed and identified relevant disclosures in both Migdal and Cosman on an element-by-element basis.

4. *Conclusions*

200. Based on my comparison of the asserted claims to the Migdal patent and Cosman, it is my opinion that the Migdal patent and Cosman render obvious claims 1-3, 5, 12-14, and 16 of the '189 patent. Furthermore, Cosman is, in fact, cited art to the Migdal patent, indicating that it was at least materially relevant to the prosecution of that patent. Moreover, both the Migdal patent and Cosman are directed to the same field and are addressing similar systems. One of ordinary skill in the art would have been motivated to combine the teachings of the Migdal patent with Cosman, for example, because both are directed to solving the problem of rapidly providing terrain texture data for displaying an image.

201. It is my understanding that the Migdal patent was cited to the Examiner during prosecution of the '189 patent. The Examiner believed that this patent disclosed all of the claims limitations except a downloading order based on resolution levels. [*See, e.g.*, GOOG 118.] There is no indication in the file history that the Examiner ever reviewed Cosman. However, Cosman discloses the very limitation that was allegedly missing from Migdal.

202. Thus, any alleged differences between claims 1-3, 5, 12-14, and 16 and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art.

V. COMPENSATION AND TESTIMONIAL EXPERIENCE

203. I am being compensated at my usual and customary rate of $600 per hour for my work in connection with this action.

204. I have testified at trial or by deposition in the following matters in the past four years:

a. 2001-2002, Macromedia Inc.: Adobe Systems Inc. v. Macromedia, Inc., Civil Action No. 00-743 JJF (United States District Court for the District of Delaware), Consultant and testifying expert witness.

b. 2003-2004, Canon Inc.: St. Clair Intellectual Property Consultants, Inc. v. Canon Inc. et al., No. 03-241 JFF (United States District Court for the District of Delaware), Consultant and testifying expert witness.

Steven K. Feiner, Ph.D.